Bounds on total economic capital: the DNB case study

Kjersti Aas[a,*], Giovanni Puccetti[b]

[a]*Department of Statistical Analysis, Image Analysis and Pattern Recognition, Norwegian Computing Center, N-0314 Oslo, Norway*
[b]*School of Economics and Management, University of Firenze, 50127 Firenze, Italy.*

Abstract

Most banks use the top-down approach to aggregate their risk types when computing total economic capital. Following this approach, marginal distributions for each risk type are first independently estimated and then merged into a joint model using a copula function. Due to lack of reliable data, banks tend to manually select the copula as well as its parameters. In this paper we assess the model risk related to the choice of a specific copula function. The aim is to compute the bounds on the total economic capital for the aggregate loss distribution of DNB, the largest Norwegian bank, and the key tool for computing these bounds is the Rearrangement Algorithm introduced in Embrechts et al. (2013). The application of this algorithm to a real situation poses a series of numerical challenges and raises a number of warnings which we illustrate and discuss.

Keywords: Model risk, Risk Aggregation, Total economic capital, Value-at-Risk, Diversification benefit, Rearrangement Algorithm.
AMS 2010 Subject Classification: 60E05, 91B30

1. Introduction

Pillar 1 (Minimum Capital Requirements) of the Basel II capital framework calculates the overall minimum capital requirement of a bank as the sum of marginal capital requirements for credit, operational and market risk. Basel II typically includes two methodologies used for computing the marginal capital charges for each of these types of risk: a standardized measurement approach and a more advanced, internal models approach.

Financial institutions are however also exposed to many other sources of risk. As cited at pages 8-9 of the consultative document Basel Committee on Banking Supervision (2010), a bank might be required to consider (in alphabetical order): asset default risk, casualty insurance risk, counterparty risk, disintermediation risk, health insurance risk, interest rate risk, lapse risk, life insurance risk, morbidity risk, mortality risk, non-life insurance risk, property insurance risk and segregated fund guarantee risk. Pillar 2 of the Basel II framework requires banks to handle all the

*Corresponding author. Department of Statistical Analysis, Image Analysis and Pattern Recognition, Norwegian Computing Center, Gaustadalléen 23, P.O. Box 114 Blindern, N-0314 Oslo, Norway.
Email addresses: Kjersti.Aas@nr.no (Kjersti Aas), giovanni.puccetti@unifi.it (Giovanni Puccetti)

relevant risk types in the so-called Internal Capital Adequacy Assessment Process (ICAAP) and a strategy for maintaining minimum capital levels. ICAAP calculations cannot result in a reduction of required capital. They will either show that additional capital is required above Pillar 1 levels or that no additional capital is needed.

Within Pillar 2 there is no regulatory guidance as to the methodology that the financial institutions should employ to integrate their different risks. Moreover, there exists no standard or state-of-the-art approach for aggregating the marginal risk types to the total risk. While the risk types to be considered and the risk metric to be used are usually set by the national authority, financial institutions are allowed to use a variety of mathematical approaches for aggregating their risks.

The task of incorporating multiple types of risks into a single metric is referred to as *Risk Aggregation*. An appropriate risk aggregation framework is fundamental for adequate firm-wide risk management. Risk aggregation models are used to support decisions about capital allocation, capital adequacy and solvency. They are also used to support risk management functions such as risk identification, monitoring and mitigation. However, risk aggregation imposes a number of significant challenges. Most importantly, a risk aggregation model should include a meaningful assumption of interdependence among the risks and it should be able to handle loss exposures that are inhomogeneous by nature and in time.

The approaches proposed in the literature for combining marginal risk information into a total economic capital can be divided into two main categories: *bottom-up* and *top-down* aggregation methods. In bottom-up aggregation, the idea is to identify the economic risk factors that are important drivers of the different risk types and develop a simultaneous model for the evolution of these risk factors. The simultaneous model includes a description of the dependence structure of the risk factors, either through a correlation matrix or a copula. The losses related to the different risk types are then determined by non-linear functions of the fluctuations in the risk factors. The final aggregation is performed by simulating a large number of scenarios (one million and more) for the risk drivers. This type of method is used for instance in Kretzschmar et al. (2010) and, partly, in Aas et al. (2007).

In the top-down aggregation approach, one develops marginal models for the yearly loss distribution of each risk type independently. These marginal distributions are then merged into a joint distribution using a correlation structure and/or a copula function. This kind of approach, which is described in Rosenberg and Schuermann (2006), Dimakos and Aas (2004) and Brockmann and Kalkbrener (2010) among others, was favored by over 75% of the surveyed banks in a comprehensive survey by the International Financial Risk Institute (IFRI Foundation and CRO Forum (2007)). Hence, in this paper we focus on a top-down aggregation approach.

Stepping from a number of marginal loss distribution models into a multivariate model for the joint risk portfolio calls for an assumption of a dependence structure among the individual random variables. This dependence structure is typically provided in the form of a copula model. Here, the reader not familiar with the theory of copulas is referred to Chapter 5 in McNeil et al. (2005). A copula function is used to merge the marginal distributions of each individual risk factor into a joint model. Since it is relatively easy to generate simulations from such a model, the final estimate of the economic capital is typically produced over a high number of simulated samples.

If one uses the top-down aggregation method to aggregate the risk types, one first has to decide

which copula to use and then estimate its parameters. Due to lack of reliable data, most banks tend to manually select the copula as well as its parameters. Inter-risk diversification is typically superimposed using rank correlations that are smaller than one, resulting in a downward adjustment to the total economic capital. The reliability of such total economic capital reductions have recently been questioned by Kretzschmar et al. (2010) and Grundke (2013).

Hence, in this paper the aim is to quantify the *model risk* related to the choice of a specific dependence structure. More specifically, we determine the best and worst possible economic capitals which are coherent with the marginals of each risk type, where the economical capital is computed in terms of the Value-at-Risk (VaR) for the total loss distribution. It is well known that the worst possible VaR does *not* correspond to maximally correlated random marginals; see Fallacy 3 in the milestone paper Embrechts et al. (2002). Based on the difference between the worst and best capital estimates and on the ratio between the worst estimate and the maximal correlation case, we will study two useful indicators of the model risk related to the choice of a specific dependence structure.

The key tool for the computation of bounds on economic capital will be the Rearrangement Algorithm (RA). The RA was introduced in Embrechts et al. (2013) for the computation of the best/worst-possible VaR for the sum of d inhomogeneous risks with given marginal distributions. While the RA is applied only to continuous, parametric families of distributions in Embrechts et al. (2013), we will handle the more realistic case for which some of the marginals are given in the form of simulated samples. The shift from continuous distributions to simulated samples poses several numerical challenges which we present and discuss. Our final aim will be the computation of the best/worst-possible 99.97%-VaR for the total loss distribution of DNB, the largest Norwegian bank. Since all underlying methods and models are adapted to the requirements in the Basel II (becoming Basel III) regulations, we however believe the results, the methodologies and guidelines described in this paper to be applicable also in a broader context.

We very much hope that our paper is both accessible to the academic researcher as well as to the more quantitative practitioner and that the more applied statements are fully acceptable to the end-user. With this goal in mind, we have kept the technical details to a minimum, stressing more the algorithmic, numerical aspects and practical consequences of the results discussed. The mathematics included is however totally correct in every respect. The reader interested in more mathematical details is always referred to the relevant research papers.

The rest of this paper is organized as follows. In Section 2 we give an overview of the methodology used by the bank to generate the marginal distributions of the different risks. Section 3 reviews the theory behind the Rearrangement Algorithm and describe its application to realistic situations like the DNB case. Finally, in Section 4 we summarize some issues and warnings concerning the computation of economic capital in a practical setting.

2. The DNB total risk model

DNB, the largest Norwegian bank, is exposed to credit, market, operational and business risk. In addition, it faces two types of risk that stems from the ownership of the group's life insurance company. The first, which here is denoted asset risk, is the risk connected to negative movements in the assets of the insurance company, while the other, denoted insurance risk, is the risk of

losses due to unforeseen increases in life claims (e.g. caused by changes in death probabilities and disability rates). In conclusion, the aggregate loss for DNB is given by the sum

$$L_6^+ = L_1 + \cdots + L_6,$$

where the random variables $L_1, \ldots, L_6$ represent marginal risk exposures for credit (L_1), market (L_2), asset (L_3), operational (L_4), business (L_5), and insurance (L_6) risks. We also denote by $F_1, \ldots, F_6$ the corresponding marginal distribution functions.

DNB has a model for computing its total economic capital that is in current use and hence plays an important role in measuring and assessing the total risk level of the financial institution. We just briefly sketch below how the six marginal models $F_1, \ldots, F_6$ are generated. For more details on the complete procedure, the interested reader is referred to Aas et al. (2007).

The credit risk distribution F_1 is generated from the widely used multi-factor normal copula model that is associated with Credit Metrics; see Goupton et al. (1997). Since DNB's credit portfolio is very large, straightforward simulation from this model would be a computationally very heavy procedure. Consequently DNB has chosen to perform the simulation in two steps: first, the loss distributions of each of the largest customers are separately simulated, and then the losses of each industry sector (without the largest obligors) are generated using an approximate formula. Since the probability of large portfolio losses is very small, naive Monte Carlo simulation as described above typically requires a very large number of runs to achieve a satisfactory variance for the VaR estimate. For the credit portfolio of DNB not even 5 million simulations are sufficient to obtain the desired accuracy for the 99.97% VaR. Hence, importance sampling is used to increase the simulation accuracy in the tail (beyond the 96% quantile). For more details on the specific procedure, see Reitan and Aas (2010).

Market risk is typically measured on a short time horizon, such as 10 days. In the DNB total risk model, however, the distributional model F_2 for market risk must be scaled to a 1-year horizon to be consistent with the other risk types. The market risk portfolio is assumed to be composed of K asset classes, where each class is assumed to depend on fluctuations of one specific risk factor, e.g. an interest rate or an exchange rate. Also connected to each class is the exposure and a liquidation period. In DNB, like in many other financial institutions, market risk is managed by applying risk limits to traders' activities. Hence, the exposure is defined as the expected utilization of the limit. Some of the asset classes are composed of both long and short positions, meaning that the net position will vary over time. To be on the conservative side, DNB assumes that they are always positioned the wrong way. To generate the market risk distribution, daily changes in the value of all risk factors are first simulated from a multivariate CCC-GARCH model (Bollerslev, 1990) with the Student's t-distribution as the conditional distribution. Then, for each asset class and all days t, the loss corresponding to day t is computed as the relative change in the risk factor during the period [t, t+liquidation period] times the exposure, and the total loss corresponding to this day is obtained summing up the losses for all asset classes. Finally, the yearly market loss is given as the maximum of the daily losses. Note that contrary to the other risk types, market risk is two-sided, meaning that one may have both gains and losses.

The modelling of the asset risk distribution (F_3) is similar to that of market risk. First, daily changes in the value of M risk factors are simulated from a multivariate CCC-GARCH model. Then, the total value of financial assets of the insurance company is computed for each day t,

taking strategies for rebalancing into account, and this time series is compared to the insurers' buffer capital to determine the loss. The yearly asset loss is finally given as the maximum of the resulting daily losses.

Due to a small database on internal operational losses, DNB currently uses the Basel II standardised approach Basel Committee on Banking Supervision (2004) for determining economic capital for operational risk. To be able to incorporate the operational risk into the same framework as the other distributions, they need however to assume a distribution for the operational losses. Based on simulation studies, the operational loss distribution F_4 is assumed to be a LogNormal. The two parameters of the LogNormal are determined assuming that the 99.9% percentile and the mode are known. The 99.9% percentile is chosen to be equal to the Basel II capital, while the mode is determined based on expert opinions. The current choice of the LogNormal distribution, and the method for estimating its parameters must be considered as preliminary. As soon as the database on internal losses is deemed to be sufficiently large, DNB will possibly replace the LogNormal assumption with a more sophisticated technique like EVT. However, this change will not affect the techniques described in the remainder of the paper.

Due to the poor quality of internally available data for business risk, and the limited size of insurance risk when compared to the other risk types, also business and insurance distributions F_5 and F_6 are simply assumed to have LogNormal distributions.

To summarize, the marginal models F_1, F_2 and F_3 are simulated samples generated via Monte Carlo simulations, while F_4, F_5 and F_6 are provided in the form of LogNormal distributions. To give an idea of the characteristics of the marginal distributions, Table 1 shows typical summary statistics (the distributions are continuously being updated). The mean values and VaR estimates are given in million Norwegian kroner (MNOK).

L_i	Risk type	Distribution function	Mean	Skewness	Kurtosis	$\text{VaR}_{0.9997}(L_i)$
L_1	Credit	simulated sample	5206.75	3.54	35.61	61912.89
L_2	Market	simulated sample	1388.93	1.30	5.70	5508.22
L_3	Asset	simulated sample	406.50	5.24	41.42	13100.73
L_4	Operational	LogNormal	840.73	3.04	20.03	7703.97
L_5	Business	LogNormal	743.34	2.12	8.94	4529.12
L_6	Insurance	LogNormal	438.98	0.77	1.09	1000.09

Table 1: Summary statistics for the different risk types. The mean values and 99.97%-quantiles are given in MNOK.

DNB computes the capital to be reserved as the Value-at-Risk (VaR) for the aggregate loss L_6^+, at the confidence level α. This VaR is is defined as

$$\mathrm{VaR}_\alpha(L_6^+) = \inf\{x \in \mathbb{R} : P(L_6^+ \leq x) > \alpha\}. \tag{2.1}$$

To aggregate its marginal risks, DNB uses the top-down approach. More specifically, a t-copula is used to model the dependence among the six risk types and the VaR is estimated via simulations. Because of the lack of a multivariate dataset for determining the appropriate copula family, DNB follows common practice and chooses a specific dependence structure among the risks entirely on the basis of *expert opinions*. In Section 3 we use the Rearrangement Algorithm to quantify the model uncertainty associated with the choice of dependence structure.

As far as the confidence level is concerned, α is set to 99.97%. This is due to the fact that DNB wants its public debt to be rated AA. The connection between the rating and confidence level may be explained as follows (Hull, 2012, Chapter 23.1): Economic capital is usually defined as the amount of capital a financial institution needs in order to absorb losses over one year with a certain confidence level. The confidence level is therefore the probability that the bank will not run out of capital (i.e., become insolvent) in the next year. A common objective for a large international bank is to maintain a AA credit rating. Corporations rated AA have a one-year probability of default of about 0.03. This suggests that the confidence level should be set as high as 99.97% for economic capital to be a guide as to what is necessary to maintain a AA rating.

3. Measuring dependence uncertainty via the Rearrangement Algorithm

As stated in Section 2, DNB uses a copula C to model the dependence among the six risk types. A copula C is a d-dimensional distribution function (df) on $[0,1]^d$ with uniform marginals. Given a copula C and d univariate marginals $F_1, \dots, F_d$, one can always define a df F on $\mathbb{R}^d$ having these marginals by

$$F(x_1, \dots, x_d) = C(F_1(x_1), \dots, F_d(x_d)),\ x_1, \dots, x_d \in \mathbb{R}. \tag{3.1}$$

Sklar's Theorem (see Theorem 5.3 in McNeil et al. (2005)) states conversely that we can always find a copula C coupling the marginals F_i of a fixed joint distribution F through the above expression (3.1). Hence Sklar's Theorem states that a copula C of a multivariate distribution F contains all the dependence information of F. In the remainder, we write $(L_1^C, \dots, L_d^C)$ to indicate a d-dimensional random vector which has fixed marginals $F_1, \dots, F_d$ and a dependence structure in the form of a copula C.

We now introduce a measure of the model risk associated to the choice of the copula on top of the marginal risk information for the six-dimensional ($d = 6$) DNB total risk portfolio. First, we define the best and worst possible VaR capitals to be held in accordance with the given marginal distributions as

$$\underline{\mathrm{VaR}}_\alpha(L_6^+) := \inf_{C \in \mathcal{C}_6} \left\{ \mathrm{VaR}_\alpha(L_1^C + \cdots + L_6^C) \right\}, \tag{3.2}$$

$$\overline{\mathrm{VaR}}_\alpha(L_6^+) := \sup_{C \in \mathcal{C}_6} \left\{ \mathrm{VaR}_\alpha(L_1^C + \cdots + L_6^C) \right\}, \tag{3.3}$$

where C_6 denotes the class of all 6-dimensional copulas. The following inequalities will automatically hold for any random vector $(L_1^C, \dots, L_6^C)$ having fixed marginals $F_1, \dots, F_6$ and arbitrary copula C:

$$\underline{\mathrm{VaR}}_\alpha(L_6^+) \leq \mathrm{VaR}_\alpha\left(L_1^C + \cdots + L_6^C\right) \leq \overline{\mathrm{VaR}}_\alpha(L_6^+).$$

The *dependence uncertainty spread* (DU-spread) of VaR is then defined as

$$\Delta\mathrm{VaR}_\alpha(L_6^+) := \overline{\mathrm{VaR}}_\alpha(L_6^+) - \underline{\mathrm{VaR}}_\alpha(L_6^+). \tag{3.4}$$

The DU-spread of VaR measures dependence model uncertainty on the capital reserve needed to offset the aggregate position L_6^+ and provides a measure of *model risk* which fits the framework set by Cont (2006). Here, the model risk derives from the uncertainty about the dependence structure coupling the estimated marginal distributions. In order to compute the DU-spread of VaR for the DNB portfolio, we use the so-called Rearrangement Algorithm.

3.1. The Rearrangement Algorithm for continuous marginal models.

The Rearrangement Algorithm (RA in the following) was introduced in Embrechts et al. (2013). The idea underlying the algorithm is very simple. The algorithm takes as input a a $(N \times d)$-matrix, and it iteratively rearrange the entries within each column of the matrix until each column is *oppositely ordered* to the sum of the other $(d-1)$ columns. Here we say that two columns $\boldsymbol{a}, \boldsymbol{b} \in \mathbb{R}^N$ are *oppositely ordered* if $(a_j - a_k)(b_j - b_k) \leq 0$ holds for all $1 \leq j, k \leq N$. Having two columns oppositely ordered means in practice that the largest component of the first is associated to the smallest component of the other and so on.

In Figure 1, we give a basic example in order to illustrate the algorithm. In (A) we start from an input 3×3 ($N = d = 3$) matrix and, as a first step, the RA rearranges its first column (indicated with an arrow) oppositely to the sum of the second and the third, which is reported using an extra column outside of the matrix. After the first iteration of the RA the matrix (B) is obtained. One can check that in (B) the first column has been rearranged so that the largest component (7) is now associated with the smallest component (1) of the sum of the other two columns. Similarly the second largest component (4) is now associated with second smallest (7), and the smallest component (1) now corresponds to the largest component (10) of the sum of the other two columns. The algorithm repeats the same procedure for the second (in (B)) and third column (in (C)) until it finds an *ordered* matrix (in (D)) in which no further changes are applicable. Depending on the initial matrix, several rearrangements of each column may be needed. However, the number of all possible different matrices which can be obtained from the input matrix by rearranging the entries in each column in a different order is finite, and this holds in general for any N and d. Since the RA finds a different matrix at each iteration, it is straightforward to state that the RA terminates in a finite number of steps; see also the proof of Theorem 2.1 in Puccetti and Rüschendorf (2012) for more details on this.

Procedure to compute $\overline{\mathrm{VaR}}_\alpha(L_6^+)$.

In order to compute $\overline{\mathrm{VaR}}_\alpha(L_6^+)$, we start from an input matrix in which the j-th column contains a N-point discretization of the part of the j-th marginal distribution F_j which is *above* the

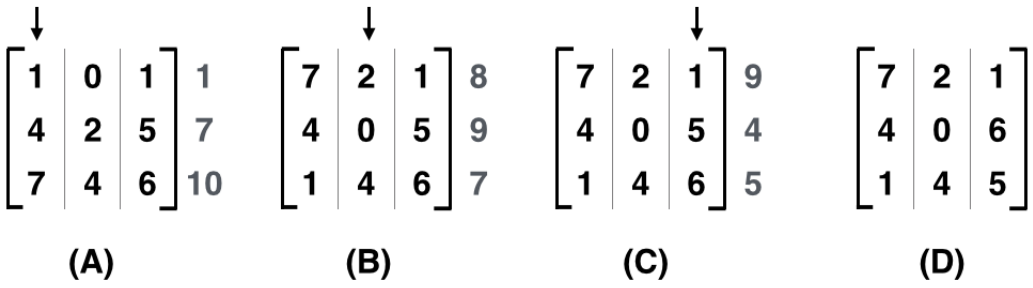

Figure 1: Iterative rearrangement of the input matrix (A) via the RA. The arrow indicates which column is rearranged at each step, while the extra column outside of the matrix is used to report the sum of the columns which are kept fixed at each step of the algorithm.

α-quantile. Then, we iteratively rearrange the entries of each column of the matrix until it stops. As a termination rule it is more convenient to require that two consecutive rearrangements of the entire set of columns provide two estimates with a fixed number of decimal digits left unchanged. At each iteration step, the *minimal* component of the vector representing the sum of all columns is taken as an estimate of $\overline{\mathrm{VaR}}_\alpha(L_6^+)$.

Rearrangement Algorithm RAU to compute $\overline{\mathbf{VaR}}_\alpha(L_6^+)$:

1. Start with an input $N \times d$ matrix $\boldsymbol{X} = (x_{i,j})$ and fix the desired level of accuracy $\epsilon > 0$.
2. Iteratively rearrange the $j-$th column of the matrix $\boldsymbol{X}$ so that it becomes oppositely ordered to the sum of the other columns, for $1 \leq j \leq d$. A matrix $\boldsymbol{Y}$ is found.
3. Repeat Step 2. until
$$s(\boldsymbol{Y}) - s(\boldsymbol{X}) < \epsilon,$$
where
$$s(\boldsymbol{X}) = \min_{1 \leq i \leq N} \sum_{1 \leq j \leq d} x_{i,j}.$$
A matrix $\boldsymbol{X}^* = \mathrm{RAU}(\boldsymbol{X})$ is found as an output of the algorithm applied to the input matrix $\boldsymbol{X}$.

If the marginal models $F_1, \ldots, F_d$ are provided in the form of continuous families of distributions, which includes the DNB risk types $4 - 5 - 6$, it is always possible to apply the RAU to two matrices, representing an approximation from below and from above of the marginal *right* α-tails. In such a way, one can obtain an interval that contains the *worst* VaR value. To this aim, define the

two matrices $\underline{\boldsymbol{X}}^{\alpha} = (\underline{x}^{\alpha}_{i,j})$ and $\overline{\boldsymbol{X}}^{\alpha} = (\overline{x}^{\alpha}_{i,j})$ as

$$\underline{x}^{\alpha}_{i,j} = F_j^{-1}\left(\alpha + \frac{(1-\alpha)(i-1)}{N}\right), \quad \overline{x}^{\alpha}_{i,j} = F_j^{-1}\left(\alpha + \frac{(1-\alpha)i}{N}\right), \tag{3.5}$$

for $1 \le i \le N, 1 \le j \le d$. Let $\underline{\boldsymbol{X}}^{*} = \mathrm{RAU}(\underline{\boldsymbol{X}}^{\alpha})$ and $\overline{\boldsymbol{X}}^{*} = \mathrm{RAU}(\overline{\boldsymbol{X}}^{\alpha})$ and define

$$\underline{s}_N = s(\underline{\boldsymbol{X}}^{*}) \quad \text{and} \quad \overline{s}_N = s(\overline{\boldsymbol{X}}^{*}).$$

Then we have $\underline{s}_N \le \overline{s}_N$ and in practice we find

$$\overline{s}_N \overset{N\to\infty}{\simeq} \underline{s}_N \overset{N\to\infty}{\simeq} \overline{\mathrm{VaR}}_{\alpha}(L_6^{+}). \tag{3.6}$$

Procedure to compute $\underline{\mathrm{VaR}}_{\alpha}(L_6^{+})$.

In order to compute $\underline{\mathrm{VaR}}_{\alpha}(L_6^{+})$, we start from an input matrix in which the j-th column contains a N-point discretization of the part of the j-th marginal distribution F_j which is *below* the α-quantile. Then, we iteratively rearrange the entries of each column of the matrix until it stops. As a termination rule again we use a condition on final accuracy considering that, at each iteration step, the *maximal* component of the vector representing the sum of all columns is taken as an estimate of $\overline{\mathrm{VaR}}_{\alpha}(L_6^{+})$.

Rearrangement Algorithm RAL to compute $\underline{\mathrm{VaR}}_{\alpha}(L_6^{+})$:

1. Start with an input $N \times d$ matrix $\boldsymbol{X} = (x_{i,j})$ and fix the desired level of accuracy $\epsilon > 0$.
2. Iteratively rearrange the $j-$th column of the matrix $\boldsymbol{X}$ so that it becomes oppositely ordered to the sum of the other columns, for $1 \le j \le d$. A matrix $\boldsymbol{Y}$ is found.
3. Repeat Step 2. until
$$t(\boldsymbol{Y}) - t(\boldsymbol{X}) < \epsilon,$$
where
$$t(\boldsymbol{X}) = \max_{1\le i\le N} \sum_{1\le j\le d} x_{i,j}.$$
A matrix $\boldsymbol{X}^{*} = \mathrm{RAL}(\boldsymbol{X})$ is found as an output of the algorithm applied to the input matrix $\boldsymbol{X}$.

Note that the RAL for the computation of $\underline{\mathrm{VaR}}_{\alpha}(L_6^{+})$ differs from the RAU for the computation of $\overline{\mathrm{VaR}}_{\alpha}(L_6^{+})$ because of the different termination conditions given in Step.3.

If the marginal models $F_1, \dots, F_d$ are provided in the form of continuous families of distributions, it is always possible to apply the RAL to two matrices, representing an approximation from below and from above of the marginal *left alpha*-parts of the distributions. In such a way, one can

obtain an interval that contains the *best* VaR value. To this aim, define the two matrices $\underline{\boldsymbol{Z}}^{\alpha} = (\underline{z}^{\alpha}_{i,j})$ and $\overline{\boldsymbol{Z}}^{\alpha} = (\overline{z}^{\alpha}_{i,j})$ as

$$\underline{z}^{\alpha}_{i,j} = F_j^{-1}\left(\frac{\alpha\,(i-1)}{N}\right), \quad \overline{z}^{\alpha}_{i,j} = F_j^{-1}\left(\frac{\alpha\, i}{N}\right),$$

for $1 \le i \le N, 1 \le j \le d$. Let $\underline{\boldsymbol{Z}}^{*} = \mathrm{RAL}(\underline{\boldsymbol{Z}}^{\alpha})$ and $\overline{\boldsymbol{Z}}^{*} = \mathrm{RAL}(\overline{\boldsymbol{Z}}^{\alpha})$ and define

$$\underline{t}_N = t(\underline{\boldsymbol{Z}}^{*}) \quad \text{and} \quad \bar{t}_N = t(\overline{\boldsymbol{Z}}^{*}).$$

Then we have $\underline{t}_N \le \bar{t}_N$ and in practice we find

$$\bar{t}_N \overset{N\to\infty}{\simeq} \underline{t}_N \overset{N\to\infty}{\simeq} \overline{\mathrm{VaR}}_{\alpha}(L_6^{+}). \tag{3.7}$$

We call the intervals $(\underline{s}_N, \bar{s}_N)$ and $(\underline{t}_N, \bar{t}_N)$ the rearrangement ranges for $\overline{\mathrm{VaR}}_{\alpha}(L_6^{+})$ and, respectively, for $\underline{\mathrm{VaR}}_{\alpha}(L_6^{+})$. For more details on the convergence of the algorithm and the limit results given in (3.6) and (3.7), we refer the reader to the remark *Convergence of the algorithm* given in Section 4.

In the case of a homogeneous portfolio of identically distributed random variables, it is possible to analytically compute $\overline{\mathrm{VaR}}_{\alpha}(L_6^{+})$ using Proposition 4 in Embrechts et al. (2013). Hence, in this case the quality of the bounds provided by the RAU can be checked. We measure the accuracy of the rearrangement range for $\overline{\mathrm{VaR}}_{\alpha}(L_6^{+})$ via the relative errors

$$\underline{e}_N = \frac{\underline{s}_N - \overline{\mathrm{VaR}}_{\alpha}(L_6^{+})}{\overline{\mathrm{VaR}}_{\alpha}(L_6^{+})} \quad \text{and} \quad \bar{e}_N = \frac{\bar{s}_N - \overline{\mathrm{VaR}}_{\alpha}(L_6^{+})}{\overline{\mathrm{VaR}}_{\alpha}(L_6^{+})}$$

and we call the interval $(\underline{e}_N, \bar{e}_N)$ the *RAU error range*.

In the homogeneous case, the analytical computation of the best VaR is possible under the assumption of an underlying marginal distribution with a decreasing density, as stated in Corollary 4.8 in Bernard et al. (2014). This assumption covers a lot of practical cases such as the Pareto distribution, but it is not satisfied in general for the LogNormal distribution. In practice, we found that the VaR lower bound stated in Bernard et al. (2014) is not sharp for any of the LogNormal densities in the DNB portfolio. Hence, in our case we are unfortunately not able to perform the same quality assessment for best VaR as for worst VaR.

However this might not be that crucial. First, the best VaR estimate relies on the N-discretization of the α *left part* of the distribution, meaning that it is reasonable to expect this estimate to be more accurate than the corresponding estimate of the worst VaR. Second, we believe the worst VaR estimate to be more interesting to a bank in general, since this is the measure that describes the worst case situation.

Since the accuracy of the worst VaR estimate can be checked only in the case of a homogeneous portfolio, we will first test the RAU on a reference marginal distribution in order to get an idea of the magnitude of the error introduced by the discretization procedure. In the case of DNB's three LogNormal distributions, the accuracy of the estimate of $\overline{\mathrm{VaR}}_{\alpha}(L_6^{+})$ seems to be decreasing with the heaviness of the right tail, as illustrated in Table 2. Hence, we choose the distribution F_4 as the

$d = 6$	$\overline{\mathrm{VaR}}_\alpha(L_6^+)$ (MNOK)	RAU error range (%)
$F_4 = \mathrm{LogN}(\mu = 6.4741049, \sigma = 0.7213475)$	56387.11	-0.0007 – 0.0004
$F_5 = \mathrm{LogN}(\mu = 6.4459970, \sigma = 0.5747400)$	31762.01	-0.0005 – 0.0003
$F_6 = \mathrm{LogN}(\mu = 6.0534428, \sigma = 0.2489544)$	6404.66	-0.0003 – 0.0001

Table 2: The worst possible 99.97%-VaR (in MNOK) and the corresponding RAU error range for the sum of $d = 6$ identically distributed random variables under different marginal assumptions. The figures in this table are computed using the RAU with $N = 10^5$ and $\epsilon = 10^{-2}$.

$d = 6$	RAU range (true=56387.1 MNOK)	RAU error range (%)
$N = 10^2$	56043.7–56661.8	-0.6090 – 0.4872
$N = 750$	56340.91–56423.05	-0.0819 – 0.0638
$N = 10^3$	56348.9–56414.8	-0.0677 – 0.0491
$N = 10^4$	56383.6–56389.8	-0.0063 – 0.0047
$N = 10^5$	56386.7–56387.4	-0.0007 – 0.0004

Table 3: The RAU range (in MNOK) and the RAU error range (in %) for the worst possible 99.97%-VaR for the sum of $d = 6$ random variables identically distributed as F_4. The true value for $\overline{\mathrm{VaR}}_\alpha(L_6^+)$ is obtained analytically using Proposition 4 in Embrechts et al. (2013).

reference marginal distribution to get a conservative estimate of the relative errors committed by the RA.

Table 3 shows the accuracy of the RAU for a portfolio consisting of $d = 6$ random variables identically distributed as F_4 for various values of N and $\epsilon = 0.01$. As shown by the table, already at $N = 10^4$, one obtains the worst VaR estimate with a relative error in the order of 10^{-5} to 10^{-4} and this in less than half a second. Unfortunately, as we will see later, the overall accuracy deteriorates when the marginal models involved are not provided in the form of continuous distributions. All the figures reported in the tables of this paper are obtained instantaneously or within seconds when computed using R on an Apple MacBook Air (2 GHz Intel Core i7, 8 GB RAM).

3.2. The RA for a combination of continuous and simulated marginal models

In Section 3.1 we defined the application of the two RAs when the marginal models are provided in the form of a parametric distributions. When simulated samples are used instead of analytically computed quantiles, the application of the two RA's are different from what is previously described:

1. *One matrix only.* For the computation of the worst VaR estimate $\overline{\mathrm{VaR}}_\alpha(L_6^+)$, the RAU uses

two matrices, representing an approximation from below and from above of the right tails of the marginal distributions under study. When the marginal distributions are provided in analytical form, these upper and lower approximations are defined in (3.5). When one has a marginal distribution provided in the form of simulated samples, it is not in general possible to obtain two discrete models approximating the marginal right tail from below and from above. In this case, the RAU is applied to a single matrix in which each column contains N observed simulations in the right $(1-\alpha)$-tail of the corresponding distribution. Formally, we define $\boldsymbol{X}^\alpha = (x^\alpha_{i,j})$, where each $(x^\alpha_{1,j}, \dots, x^\alpha_{N,j})$ is a vector of N simulated samples in $[F_j^{-1}(\alpha), F_j^{-1}(1)]$, $1 \le j \le d$. Using the notation defined in Section 3.1, one applies the RAU to the matrix $\boldsymbol{X}^\alpha$ and computes $s(\mathrm{RAU}(\boldsymbol{X}^\alpha))$ as an estimate of $\overline{\mathrm{VaR}}_\alpha(L_6^+)$.

Analogous considerations hold for the computation of $\underline{\mathrm{VaR}}_\alpha(L_6^+)$. In this case, the RAL is applied to a single matrix in which each column contains N observed simulations in the left α-tail of the corresponding distribution. Formally, we define $\boldsymbol{Z}^\alpha = (z^\alpha_{i,j})$ where each $(z^\alpha_{1,j}, \dots, z^\alpha_{N,j})$ is a vector of N simulated samples in $[F_j^{-1}(0), F_j^{-1}(\alpha)]$, $1 \le j \le d$. Then, one applies the RAL to the matrix $\boldsymbol{Z}^\alpha$ and computes $t(\mathrm{RAL}(\boldsymbol{Z}^\alpha))$ as an estimate of $\underline{\mathrm{VaR}}_\alpha(L_6^+)$.

2. *Random estimates.* As one relies on a random sample, two different set of simulations produce two different estimates for the VaR. No deterministic ranges can be obtained. However, empirical confidence intervals for the random estimate can be obtained by running the algorithms a large number of times.

3. *Choice of N.* The constant N denotes the number of simulated samples in the right or left tail of each of the marginal models. Assume that the models used by a bank are able to produce a number of M simulations over the entire domain of each of the marginal distributions F_j. This means that the RAU can be applied to a matrix with at most $N = M(1-\alpha)$ rows, while the RAL can be applied to a matrix with at most $N = M\alpha$ rows. This point is particularly relevant. With a probability level of $\alpha = 99.97\%$, if one is able to obtain $M = 10^6$ simulations, only $N = 300$ of them will be used in the RAU, while $N = 999700$ of them will be used in the RAL. Hence, for a given M, we expect the RAL to be more accurate than the RAU.

Now we test the accuracy of RAU in a homogeneous portfolio where the marginals are provided partly in the form of continuous models and partly in the form of simulated samples. To this aim, we apply the RAU to a $N \times 6$ matrix where columns 1–3 contain simulated samples in the right $(1-\alpha)$ tail of F_4, while columns 4–6 contain the deterministic quantiles of F_4. Hence, we apply the RAU to the matrix $\boldsymbol{X}^\alpha = (x^\alpha_{i,j})$, where

$$x^\alpha_{i,j} = s_{i,j},\ 1 \le j \le 3, \quad \text{and} \quad x^\alpha_{i,j} = F_4^{-1}\left(\alpha + \frac{(1-\alpha)i}{N}\right),\ 4 \le j \le 6, \tag{3.8}$$

for $1 \le i \le N$. Here each $(s_{1,j}, \dots, s_{N,j})$, $1 \le j \le 3$, is a vector of N simulated observations in $[F_4^{-1}(\alpha), F_4^{-1}(1)]$. Contrary to the definition of the input matrices in (3.5), notice that the first three columns of the input matrix $\boldsymbol{X}^\alpha$ in (3.8) are not ordered; see also the remark *To randomize or not to randomize?* given in Section 4 below.

At this point, we run the RAU 10^5 times in order to compute empirical confidence intervals for $\overline{\mathrm{VaR}}_\alpha$. The results are shown in Table 4. Although the accuracy of the RA, measured in terms of 95% empirical confidence interval for the relative error, is worse than that shown in Table 3, the overall performance of the algorithm is still more than reasonable.

It is important to notice that the relative errors reported in Table 4 are only valid when the random variables are identically distributed as F_4 and are only to be used as guidelines for the choice of N. We stress again that, at the moment, there is no way one could check the accuracy of the RA for an inhomogeneous portfolio of risks with (some of the) marginals provided in the form of simulated samples. However, since the reference distribution used is the one producing the poorest accuracy performance (see Table 2) for an homogeneous portfolio, we expect the final relative errors committed in the real DNB case not to be larger than the ones reported in Table 4.

$d = 6$	RAU 95% CI (true=56387.1)	RA error 95% CI (%)
$N = 10^2$	55891.98–57167.62	-0.8781 – 1.3842
$N = 750$	56171.67–56643.03	-0.3821 – 0.4539
$N = 10^3$	56192.52–56609.12	-0.3451 – 0.3937
$N = 10^4$	56321.69–56452.96	-0.1160 – 0.1168
$N = 10^5$	56366.46–56407.96	-0.0366 – 0.0370

Table 4: 95% confidence intervals (CI) for the RAU estimate and RAU error for the worst possible 99.97%-VaR for the sum of $d = 6$ random variables identically distributed as F_4, when three of the marginals are provided in the form of simulated samples. The empirical confidence intervals are evaluated over 10^5 runs of the algorithm. The true value for $\overline{\mathrm{VaR}}_\alpha(L_6^+)$ is obtained analytically using Proposition 4 in Embrechts et al. (2013).

3.3. The RA for the DNB case

Due to memory constraints, the DNB model is able to generate at most $M = 2.5 \times 10^6$ simulations for credit, market and asset risk. Using 2.5 million simulations may seem sufficient for computing the 99.97% quantile. However, this means that the RAU can be applied to an input matrix with only $N = (1 - \alpha)M = 750$ rows. On the other hand, the RAL can be applied to an input matrix with $N = \alpha \times M = 2499250$ rows, and it is reasonable to expect the best VaR estimate to be more accurate than the worst VaR estimate. For operational, business and insurance risk we used deterministic quantiles following the technique described in Section 3.2.

In Table 5, we report the best and worst possible VaR estimates obtained for the DNB model. Table 5 also shows the DU-spread of VaR_α as defined in (3.4) and the *comonotonic* VaR which corresponds to the case of maximally correlated marginal exposures. We have also compared the comonotonic VaR value with the *independence* VaR value obtained under the assumption of independence and the *t-copula* value resulting when a t-copula with 6 degrees of freedom is used to model the dependence among the six risk types (which is the model used in practice by DNB). The

comonotonic VaR is simply obtained by summing up the marginal quantiles in Table 1 (see McNeil et al. (2005, Proposition 6.15)), while the independence and t-copula values are obtained via simulations. As we already remarked, for $\overline{\mathrm{VaR}}_\alpha$ one can reasonably assume that the relative error committed by the RAU is not larger than the one obtained in Table 4, while we expect $\underline{\mathrm{VaR}}_\alpha$ to be much more accurate. The DU-spread of the VaR reported in Table 5 can be used as a measure of model risk related to the uncertainty about the structure of dependence amongst the marginals. Banks typically have better methods for estimating the marginal distributions for each risk type than they have for the dependency structure. Hence, it is very useful for a bank to know the best and especially the worst possible VaR for a given a set of marginals. Once one is able to quantify DU-spreads, the following regulatory task would probably be to assess a corresponding extra capital to be reserved to offset model risk. This a delicate point that goes well beyond the scope of our paper. Hence, we deliberately want to leave this issue open to encourage further discussions within banking practice and regulation.

$\underline{\mathrm{VaR}}_\alpha(L_6^+)$	independence	t-6 copula	comonotonic	$\overline{\mathrm{VaR}}_\alpha(L_6^+)$	$\Delta\mathrm{VaR}_\alpha(L_6^+)$
62156.4	66182.84	77896.33	93755.02	105878.2	43721.8

Table 5: From left to right: best, independence, t-6 copula, comonotonic, and worst VaR estimates and the DU-spread of the VaR for the DNB case. All numbers are given in MNOK and the quantile level is set to $\alpha = 99.97\%$.

Diversification ratio and diversification benefit

Define the *diversification ratio* (SR) for the aggregate loss $L_6^+ = \sum_{i=1}^6 L_i$ as

$$\triangle_\alpha(L_6^+) = \frac{\mathrm{VaR}_\alpha(L_6^+)}{\sum_{i=1}^6 \mathrm{VaR}_\alpha(L_i)}. \tag{3.9}$$

The diversification ratio measures the ratio between the total VaR of a portfolio and the sum of the marginal exposures. Superadditivity of VaR can be also described using the equivalent concept of *diversification benefit* formally introduced in Cope et al. (2009), but the ratio

$$\mathrm{VaR}_\alpha(L_6^+)/(\sum_{i=1}^6 \mathrm{VaR}_\alpha(L_i))$$

was already studied in the milestone paper by Embrechts et al. (2002).

A diversification ratio between 0 and 1 indicates that diversification effects occur in the portfolio, meaning that the aggregate position L_6^+ is less risky than the sum of the marginal exposures. This is the typical situation occurring for elliptically distributed risk portfolios, for which it is well known that $\triangle_\alpha(L_6^+) \in [0, 1]$; see Theorem 6.8 in McNeil et al. (2005).

The case $\triangle_\alpha(L_6^+) = 1$ occurs when the risks $L_1, \ldots, L_6$ are *comonotonic*, i.e. they are all almost surely increasing functions of a common random factor. In this case it is well known that the VaR for their sum is equal to the sum of marginal VaR numbers; see Proposition 6.15 in McNeil et al. (2005). See also Section 6.2.2 in McNeil et al. (2005) for a detailed discussion of the concept of

comonotonicity within quantitative risk management. The comonotonic case represents maximal correlation between the marginals and it is erroneously considered as highly conservative. As shown by Table 5, the worst-case total risk of DNB exceeds the sum of the marginal risks, meaning that the value of $\triangle_\alpha(L_6^+)$ may well be larger than 1. We define the *worst diversification ratio* (WSR) for L_6^+ as

$$\overline{\triangle}_\alpha(L_6^+) = \frac{\overline{\mathrm{VaR}}_\alpha(L_6^+)}{\sum_{i=1}^6 \mathrm{VaR}_\alpha(L_i)}.$$

For DNB at the quantile level $\alpha = 99.97\%$, we have $\overline{\triangle}_\alpha(L_6^+) = 1.1293$. It is well known that the WSR may be larger than one, especially in the case of infinite mean models; see for instance Ibragimov and Walden (2008) and Mainik and Rüschendorf (2010). In the DNB example however, the marginal risks possess a finite expectation.

4. Remarks and warnings

We conclude our paper by discussing some important issues that the assessment of model uncertainty in a real situation raises.

Only tails matter. As a consequence of Theorem 2.1 in Puccetti and Rüschendorf (2013), the worst-possible α-VaR for a sum of random variables only depends on the right $(1-\alpha)$ tails of the support of the fixed marginal distributions. Maximization of the VaR of L_6^+ can be equivalently seen as the maximization of the survival function $P(L_6^+ \geq s)$, for some real threshold s. Since the probability $P(L_6^+ \geq s)$ is increasing on the mass of the right tail of the fixed marginals and no more than $(1-\alpha)$ probability mass can be allocated to the right tail of the optimal solution, it is intuitively obvious that the optimal solution should use only the largest $(1-\alpha)$ part of each marginal component. As a result, the optimal dependence structure is to be determined only in the right tails of the marginal supports, while interdependence in the lower parts can be set arbitrarily. In order to apply the RAU for the worst VaR, a financial institution has to be primarily (if not totally) interested in providing a robust statistical model or accurate simulations for the right tails of its marginal risks. However, there are also several reasons why a financial institution would want to model/simulate the whole distribution and not only the tail. For example, DNB uses the total risk model for other tasks than economic capital calculation. When e.g. computing earnings-at risk and the best VaR estimate $\underline{\mathrm{VaR}}_\alpha(L_6^+)$, the lower quantiles are of relevance.

Uncorrelation is not meaningful. An implication of the previous statement is that it is always possible to find a worst-case VaR copula having any dependence measure like standard correlation or Kendall's/Spearman's rank correlation equal to zero. Thus, assuming that overall correlation between the risks of the bank is small does not have any direct consequence on the worst VaR estimate. Moreover, contrary to common belief, assuming the marginal risks to be positively dependent, has a minimal impact on the estimate. See the discussion carried out in Section 3 in Embrechts et al. (2013).

More information on the dependence structure. Providing information about dependence within subgroups of the random variables will considerably lower the worst VaR estimate. The RAU and RAL algorithms can also be applied in the case where one is able to estimate the copula for some bivariate pairs of random variables; see Section 4 in Embrechts et al. (2013). Recently, Bernard et al. (2014b) introduced the Extended Rearrangement Algorithm (ERA) that allows to approximate sharp VaR bounds on a sum of risks when the variance of the portfolio sum is also known. Adding the variance constraint gives rise to significantly tighter bounds in all situations of interest.

Copulas corresponding to worst-possible VaR scenarios. Any output matrix from the RAs can be seen as the support of a N-discrete, d-variate distribution giving probability mass $1/N$ to each one of its N row vectors. Thus, the algorithm described in this paper can also be used to obtain a discrete image of the structure of dependence (i.e. of the copula) attaining the worst/best possible VaR estimates. We call any such copula a *worst-possible* or *best-possible* VaR *scenario*. In Figure 2 we give bivariate projections of the upper-orthant part of the support of the worst-possible VaR scenario in the case of the sum of $d = 3$ random losses identically distributed as F_4. The number of points used in the discretization of the copula is set to $N = 750$ to resemble the real application discussed in Section 3.2. The figure clearly shows the optimal dependence structure underlying a worst-possible VaR scenario. This scenario is *not* the one for which all marginals are positively dependent, i.e. where they all tend to be large at the same time. As already remarked, maximization of the VaR of a sum is equivalent to the maximization of the tail function $P(L_6^+ \geq s)$. An optimal allocation of the $(1 - \alpha)$ probability mass is obtained either if at most one margin is exceeding the threshold s, or if the sum of all margins is exactly adding up to s. In the latter case, the variance of their sum is equal to zero, and we have a so-called *completely mixable* behavior. The interested reader can compare Figure 2 with Figures 2-3 in Embrechts and Höing (2006) and Figure 4 in Puccetti and Wang (2014). For further details on the concept of complete mixability, we refer the reader to the papers Wang and Wang (2011), Puccetti et al. (2012). Similar dependence structures arise as the solution of VaR maximization/minimization problems possibly subject to different constraints; see for example Bernard et al. (2013) and Bernard et al. (2014b). In higher dimensions, the behavior of a homogeneous portfolio is similar, but less evident, because the completely mixable region gets larger with increasing dimensions. In Figure 3 we give bivariate projections of a worst-possible VaR scenario in the case of the sum of $d = 6$ random losses identically distributed as F_4. Worst/best VaR scenarios may also be modeled by *patchwork copulas*; see Durante et al. (2013) on this.

Figure 4 shows the projections of the upper-orthant part of the support of the worst-possible VaR scenario in the case a portfolio of $d = 3$ risks having marginal distributions F_4, F_5, F_6. The cloud of points visible in each sub-figure corresponds to the region of minimal variance for the sum. For general (possibly inhomogeneous) marginals, VaR is then still maximized by copulas containing *jointly mixable* (i.e. minimal variance) regions in their upper-orthant parts. The concept of joint mixability has been introduced in the paper Wang et al. (2013). Analogous patterns are found in Figure 5, which shows the projections of the upper-right part of the support of the worst-possible VaR scenario for the six-dimensional DNB case.

Value-at-Risk versus Expected Shortfall. In the DNB case, where the marginal exposures possess a finite expectation, Expected Shortfall (ES) is a natural alternative to VaR.

For a random variable X with $\mathbb{E}[|X|] < \infty$ and distribution function F_X, the expected shortfall (ES) at confidence level $\alpha \in (0, 1)$ is defined as

$$\mathrm{ES}_\alpha(X) = \frac{1}{1-\alpha} \int_\alpha^1 \mathrm{VaR}_q(X)\, dq. \tag{4.1}$$

Unlike VaR, ES accounts for risk in a more comprehensive manner considering both the size and likelihood of losses above a certain threshold (e.g. the 99.97%-quantile). Most importantly, ES is a subadditive risk measure and it is well known that the worst-possible value in the case of an aggregate position L_6^+ is attained when the risks are comonotonic, i.e. perfectly positively dependent. Formally stated, one has that

$$\overline{\mathrm{ES}}_\alpha(L_6^+) := \sup_{C \in \mathcal{C}_6} \left\{ \mathrm{ES}_\alpha(L_1^C + \cdots + L_6^C) \right\} = \mathrm{ES}_\alpha(L_1^M + \cdots + L_6^M) = \mathrm{ES}_\alpha(L_1) + \cdots + \mathrm{ES}_\alpha(L_6),$$

where M denotes the comonotonic copula; see for instance Section 1 in Puccetti (2013).

A model risk analysis like the one previously carried out for VaR is also possible using ES as the benchmark risk measure. Analogously to (3.2) and (3.4), one can define the best-possible ES capital coherently with the fixed marginal distributions $F_1, \ldots, F_d$ as

$$\underline{\mathrm{ES}}_\alpha(L_6^+) := \inf_{C \in \mathcal{C}_6} \left\{ \mathrm{ES}_\alpha(L_1^C + \cdots + L_6^C) \right\},$$

and the DU-spread of ES_α as

$$\Delta\mathrm{ES}_\alpha(L_6^+) = \overline{\mathrm{ES}}_\alpha(L_6^+) - \underline{\mathrm{ES}}_\alpha(L_6^+).$$

The DU-spread of ES_α for the DNB portfolio can be calculated using the RA following the technique described in Puccetti (2013). The results are given in Table 6. For the sake of comparison we also repeat the VaR numbers from Table 5.

$\alpha = 99.97\%$	$\underline{\mathrm{VaR}}_\alpha(L_6^+)$	$\overline{\mathrm{VaR}}_\alpha(L_6^+)$	$\Delta\mathrm{VaR}_\alpha(L_6^+)$
	62156.4	105878.2	43721.8
$\alpha = 99.97\%$	$\underline{\mathrm{ES}}_\alpha(L_6^+)$	$\overline{\mathrm{ES}}_\alpha(L_6^+)$	$\Delta\mathrm{ES}_\alpha(L_6^+)$
	74354.7	110588.8	36234.1

Table 6: Upper row from left to right: Best VaR estimate, worst VaR estimate and the DU-spread of VaR for the DNB case. Lower row from left to right: Best ES estimate, worst ES estimate and the DU-spread of ES for the DNB case. All numbers are given in MNOK.

By definition, ES is a more conservative risk measure than VaR, i.e. $\mathrm{ES}_\alpha(L_6^+) \geq \mathrm{VaR}_\alpha(L_6^+)$. The same inequality obviously holds also for the corresponding best/worst cases, i.e. $\overline{\mathrm{ES}}_\alpha(L_6^+) \geq$

$\overline{\mathrm{VaR}}_\alpha(L_6^+)$ and $\underline{\mathrm{ES}}_\alpha(L_6^+) \geq \underline{\mathrm{VaR}}_\alpha(L_6^+)$. As shown by Table 6, for the DNB case, the worst-possible ES estimate is just 4% bigger than the corresponding VaR estimate. This very small difference is not surprising, since conservative VaR- and ES-based capital charges have recently been shown to be asymptotically equivalent. In practice, under very weak conditions on the marginal distributions of a risk portfolio we have that

$$\lim_{d\to\infty} \frac{\overline{\mathrm{ES}}_\alpha(L_d^+)}{\overline{\mathrm{VaR}}_\alpha(L_d^+)} = 1,$$

where $L_d^+ = L_1 + \cdots + L_d$; see Embrechts et al. (2014) for a proof and the history of this limit result. Concerning Table 6, it is also interesting to mention that ES typically has a smaller DU-spread if compared to VaR. As shown and discussed in Embrechts et al. (2014), this conclusion holds asymptotically for all $\mathrm{ES}_q(L_6^+)$ vs $\mathrm{VaR}_\alpha(L_6^+)$, $q \leq \alpha$.

Convergence of the algorithm. Unfortunately, there does not exist an analytic proof that the limit results in (3.6) and (3.7) hold for all initial configurations of the algorithm. In general we have that $\underline{s}_N \leq \overline{\mathrm{VaR}}_\alpha(L_6^+)$ and $\bar{t}_N \geq \underline{\mathrm{VaR}}_\alpha(L_6^+)$ for sufficiently large N. However, the inequalities $\bar{s}_N \geq \overline{\mathrm{VaR}}_\alpha(L_6^+)$ and $\underline{t}_N \leq \underline{\mathrm{VaR}}_\alpha(L_6^+)$ are not always satisfied. This is not in contrast with convergence of the algorithm, but only means that these sequences might not be monotonic.

It is quite easy to build an example with marginals uniformly distributed in the unit interval in which the two sequences $\underline{s}_N$ and $\bar{s}_N$ do not converge to $\overline{\mathrm{VaR}}_\alpha(L_6^+)$ and the two sequences $\underline{t}_N$ and $\bar{t}_N$ do not converge to $\underline{\mathrm{VaR}}_\alpha(L_6^+)$. These examples are however based upon a special choice of the starting matrix of the algorithm. In the applications given in this paper, where the starting matrix contains ordered columns and/or simulated values, we always found the algorithm to provide excellent approximations even with moderately large values of N.

The convergence of different versions of the RA has been tested in a variety of applications to portfolios of interest in quantitative risk management; see e.g. the paper Bernard et al. (2014b) However, an analytical proof of the convergence of the RA remains an open problem.

To randomize or not to randomize? The input matrix used by the RAs in Section 3 is a *comonotonic* matrix having all its columns arranged in increasing order. In the original version of the algorithm described in Embrechts et al. (2013), the authors use a *randomized* starting matrix in order to minimize the probability of choosing a starting point not leading to convergence of the algorithm. We tried different deterministic/random choices for the starting rearrangement, but we always found the comonotonic matrix to yield the same accuracy as the randomized one. Hence, the choice of the starting matrix does not seem to be crucial in the application of the RA with a large N. The advantage of using a comonotonic starting matrix is that the VaR ranges reported in Table 3 are deterministic and do not vary across multiple repetitions of the algorithm. On the other hand, this starting matrix may give negative values for the upper error $\bar{e}_N$.

Acknowledgements

Both authors are grateful to Paul Embrechts and three anonymous referees for their very careful reading of an earlier version of this paper and various helpful comments. Giovanni Puccetti

acknowledges fruitful discussions with Carole Bernard, Paul Embrechts and Ludger Rüschendorf about the choice of the starting matrix for the RA.

The work conducted by Kjersti Aas is sponsored by Statistics for Innovation, (sfi)^2. Giovanni Puccetti acknowledges grant under the call PRIN 2010-2011 from MIUR within the project *Robust decision making in markets and organisations*.

Aas, K., X. K. Dimakos, and A. Øksendal (2007). Risk capital aggregation. *Risk Management 9*, 82–107.

Basel Committee on Banking Supervision (2004). *International Convergence of Capital Measurement and Capital Standards*. Basel: Bank for International Settlements.

Basel Committee on Banking Supervision (2010). *Developments in modelling risk aggregation.* Basel: Bank for International Settlements.

Bernard, C., X. Jiang, and R. Wang (2014). Risk aggregation with dependence uncertainty. *Insurance Math. Econom. 54*, 93 – 108.

Bernard, C., Y. Liu, N. MacGillivray, and J. Zhang (2013). Bounds on capital requirements for bivariate risk with given marginals and partial information on the dependence. *Depend. Model. 1*, 37–53.

Bernard, C., L. Rüschendorf, and S. Vanduffel (2014b). Value-at-risk bounds with variance constraints. Preprint, available at SSRN: http://ssrn.com/abstract=2342068.

Bollerslev, T. (1990). Modelling the coherence in short-run nominal exchange rates: A multivariate generalized ARCH model. *Rev. Econ. Stat. 72*, 498–505.

Brockmann, M. and M. Kalkbrener (2010). On the aggregation of risk. *J. Risk 12*, 45–68.

Cont, R. (2006). Model uncertainty and its impact on the pricing of derivative instruments. *Math. Financ. 16*(3), 519 – 547.

Cope, E., G. Mignola, G. Antonini, and R. Ugoccioni (2009). Challenges and pitfalls in measuring operational risk from loss data. *J. Operational Risk 4*(4), 3–27.

Dimakos, X. K. and K. Aas (2004). Integrated risk modelling. *Stat. Model. 4*, 265–277.

Durante, F., J. Fernández-Sánchez, and C. Sempi (2013). Multivariate patchwork copulas: a unified approach with applications to partial comonotonicity. *Insurance Math. Econom. 53*(3), 897–905.

Embrechts, P. and A. Höing (2006). Extreme VaR scenarios in higher dimensions. *Extremes 9*, 177–192.

Embrechts, P., A. J. McNeil, and D. Straumann (2002). Correlation and dependence in risk management: properties and pitfalls. In *Risk management: Value at Risk and Beyond*, pp. 176–223. Cambridge: Cambridge Univ. Press.

Embrechts, P., G. Puccetti, and L. Rüschendorf (2013). Model uncertainty and VaR aggregation. *J. Bank. Financ. 37*(8), 2750 – 2764.

Embrechts, P., B. Wang, and R. Wang (2014). Aggregation-robustness and model uncertainty of regulatory risk measures. Preprint, ETH Zurich.

Goupton, G. M., C. C. Finger, and M. Bhatia (1997). Creditmetrics. Technical document.

Grundke, P. (2013). On the reliability of integrated risk measurement in practice. *J. Risk 15*(3), 87–110.

Hull, J. (2012). *Risk Management and Financial Institutions*. 3rd edition, Wiley, Hoboken, New Jersey.

Ibragimov, R. and J. Walden (2008). Portfolio diversification under local and moderate deviations from power laws. *Insurance Math. Econom. 42*(2), 594–599.

IFRI Foundation and CRO Forum (2007). Insights from the joint IFRI/CRO forum survey on economic capital practice and applications. KPMG Business Advisory Services.

Kretzschmar, G., A. J. McNeil, and A. Kirchner (2010). Integrated models of capital adequacy – why banks are undercapitalised. *J. Bank. Financ. 34*(12), 2838 – 2850.

Mainik, G. and L. Rüschendorf (2010). On optimal portfolio diversification with respect to extreme risks. *Finance Stoch. 14*(4), 593–623.

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

Puccetti, G. (2013). Sharp bounds on the expected shortfall for a sum of dependent random variables. *Stat. Probabil. Lett. 83*(4), 1227–1232.

Puccetti, G. and L. Rüschendorf (2012). Computation of sharp bounds on the distribution of a function of dependent risks. *J. Comput. Appl. Math. 236*(7), 1833–1840.

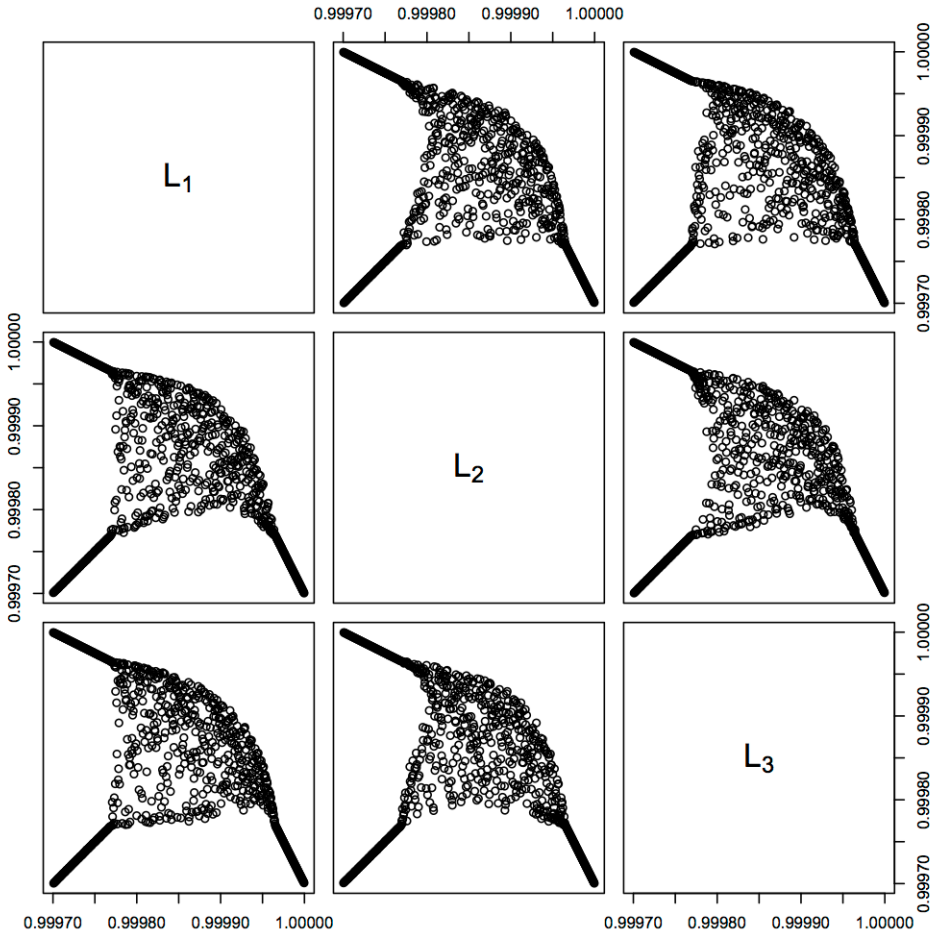

Figure 2: Two-dimensional projections to $[0.9997, 1]^2$ of the support of the discrete image of the copula attaining the worst-possible VaR for the sum of $d = 3$ random variables identically distributed as F_4.

Puccetti, G. and L. Rüschendorf (2013). Sharp bounds for sums of dependent risks. *J. Appl. Probab. 50*(1), 42–53.

Puccetti, G., B. Wang, and R. Wang (2012). Advances in complete mixability. *J. Appl. Probab. 49*(2), 430–440.

Puccetti, G. and R. Wang (2014). General extremal dependence concepts. Preprint, available at SSRN: http://ssrn.com/abstract=2436392.

Reitan, T. and K. Aas (2010). A new robust importance-sampling method for measuring value-at-risk and expected shortfall allocation for credit portfolios. *J. Credit Risk 6*(4), 1–37.

Rosenberg, J. V. and T. Schuermann (2006). A general approach to integrated risk management with skewed, fat-tailed risks. *J. Financ. Econ. 79*(3), 569 – 614.

Wang, B. and R. Wang (2011). The complete mixability and convex minimization problems with monotone marginal densities. *J. Multivariate Anal. 102*(10), 1344–1360.

Wang, R., L. Peng, and J. Yang (2013). Bounds for the sum of dependent risks and worst value-at-risk with monotone marginal densities. *Finance Stoch. 17*(2), 395–417.

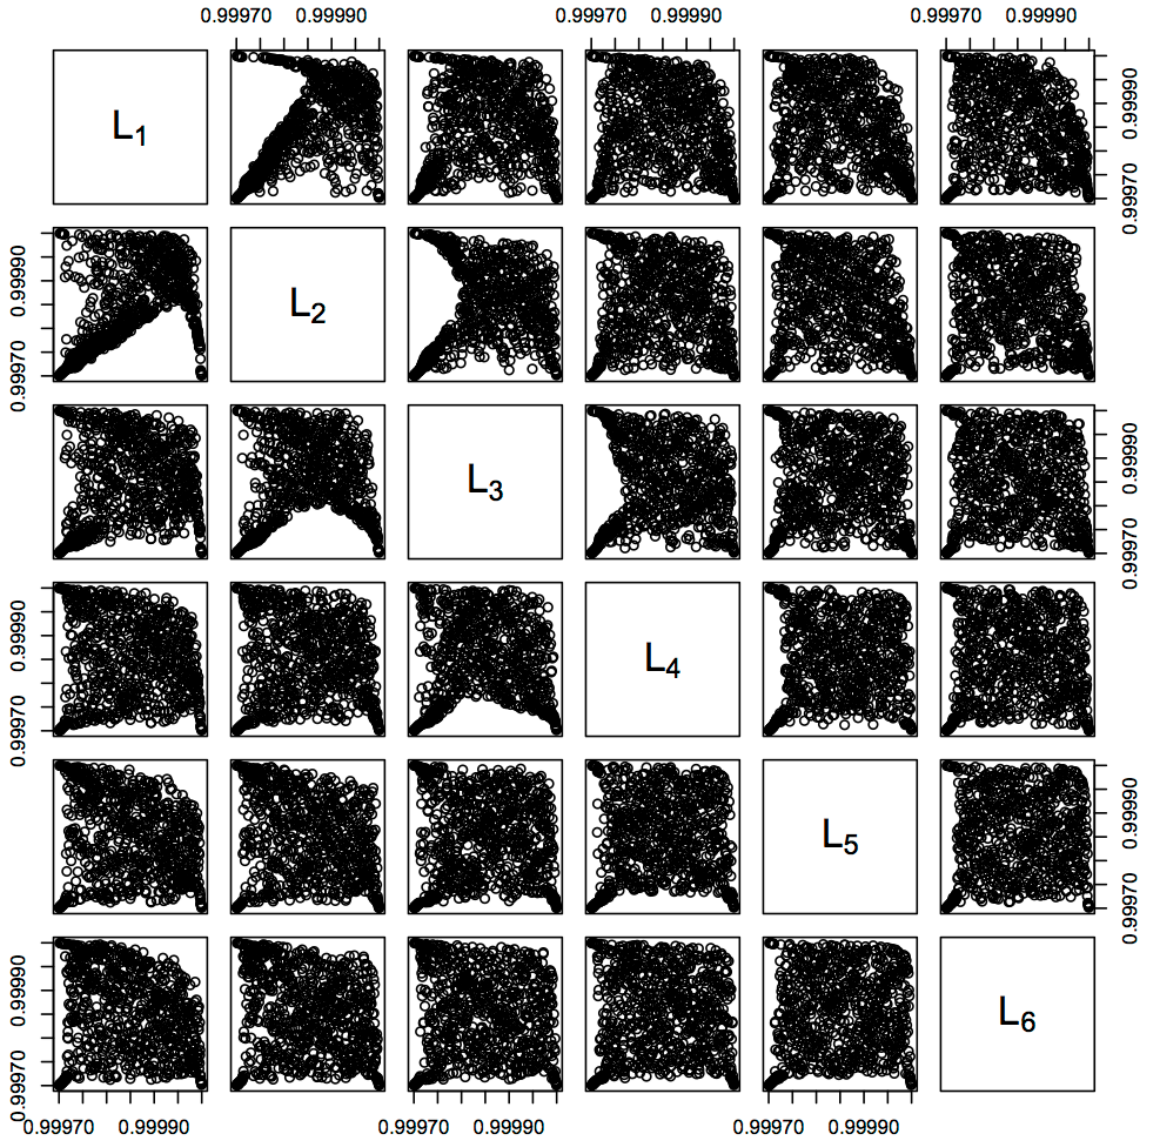

Figure 3: Two-dimensional projections to $[0.9997, 1]^2$ of the support of the discrete image of the copula attaining the worst-possible VaR for the sum of $d = 6$ random variables identically distributed as F_4.

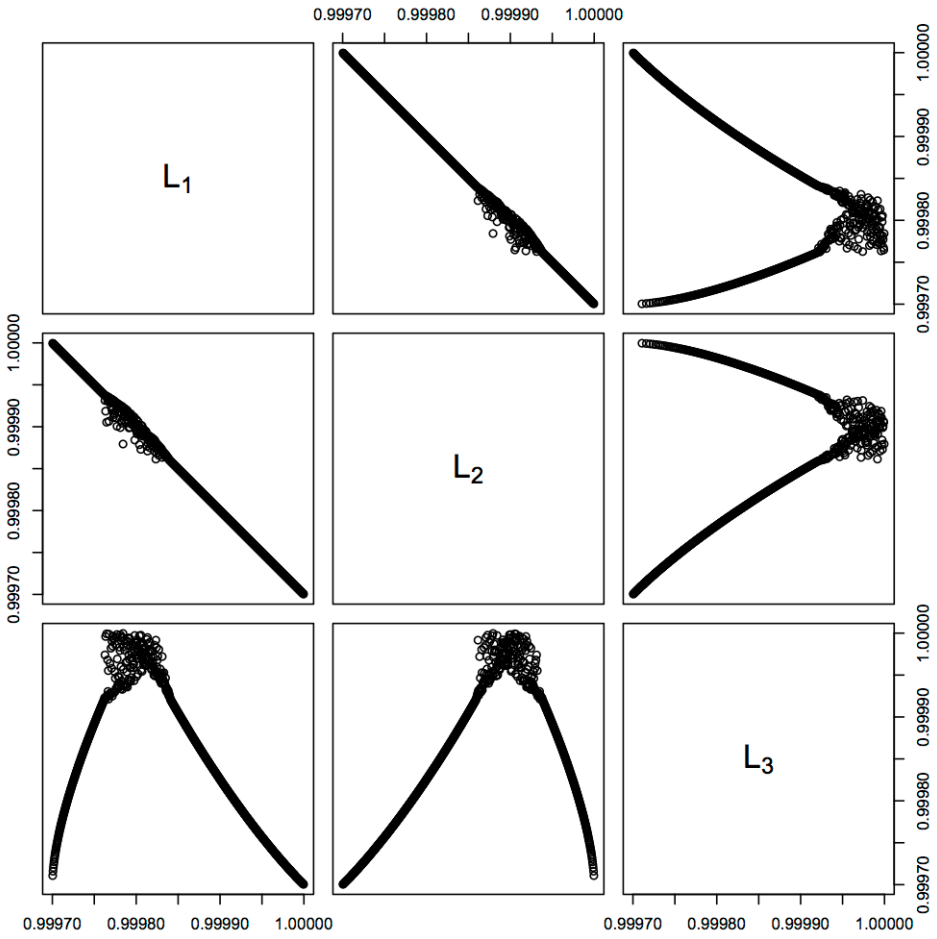

Figure 4: Two-dimensional projections to $[0.9997, 1]^2$ of the support of the discrete image of the copula attaining the worst-possible VaR for the sum of $d = 3$ random variables having marginals F_4, F_5, F_6.

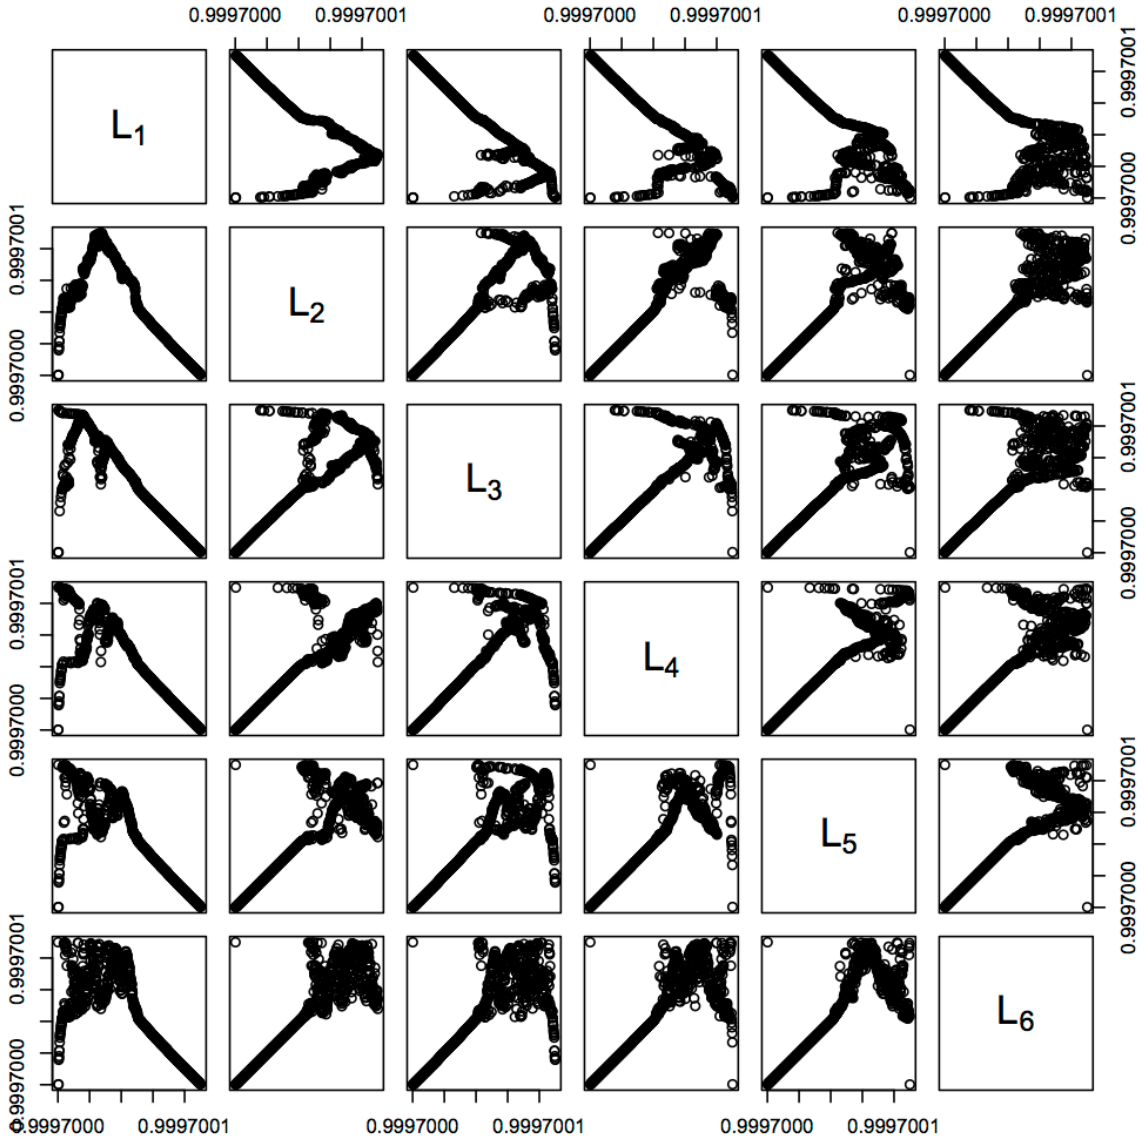

Figure 5: Two-dimensional projections to $[0.9997, 1]^2$ of the support of the discrete image of the copula attaining the worst-possible VaR for the DNB case.