



HAL
open science

Fast remote but not extreme quantiles with multiple factors. Applications to Solvency II and Enterprise Risk Management

Matthieu Chauvigny, Laurent Devineau, Stéphane Loisel, Véronique Maume-Deschamps

► To cite this version:

Matthieu Chauvigny, Laurent Devineau, Stéphane Loisel, Véronique Maume-Deschamps. Fast remote but not extreme quantiles with multiple factors. Applications to Solvency II and Enterprise Risk Management. *European Actuarial Journal*, 2011, 1 (1), pp.131-157. hal-00517766v2

HAL Id: hal-00517766

<https://hal.science/hal-00517766v2>

Submitted on 13 Jun 2011

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Fast remote but not extreme quantiles with multiple factors. Applications to Solvency II and Enterprise Risk Management

Matthieu Chauvigny

Milliman, Paris, 64 rue Pierre Charron, 75008 Paris, email:
Matthieu.Chauvigny@Milliman.com

Laurent Devineau

Milliman, Paris, 64 rue Pierre Charron, 75008 Paris, email:
laurent.devineau@milliman.com

Université de Lyon, Université Lyon 1, Institut de Science Financière et d'Assurances,
Laboratoire SAF EA 2419, 50 Avenue Tony Garnier, F-69007 Lyon, France

Stéphane Loisel

Université de Lyon, Université Lyon 1, Institut de Science Financière et d'Assurances,
Laboratoire SAF EA 2419, 50 Avenue Tony Garnier, F-69007 Lyon, France
Ph.: +33437287429, Fax: +33437287632, email: stephane.loisel@univ-lyon1.fr

Véronique Maume-Deschamps

Université de Lyon, Université Lyon 1, Institut de Science Financière et d'Assurances,
Laboratoire SAF EA 2419, 50 Avenue Tony Garnier, F-69007 Lyon, France
Ph.: +33437287636, Fax: +33437287632, email: veronique.maume@univ-lyon1.fr

ABSTRACT

For operational purposes, in Enterprise Risk Management or in insurance for example, it may be important to estimate remote (but not extreme) quantiles of some function f of some random vector. The call to f may be time- and resource-consuming so that one aims at reducing as much as possible the number of calls to f . In this paper, we propose some ways to address this problem of general interest. We then numerically analyze the performance of the method on insurance and Enterprise Risk Management real-world case studies.

Keywords: Quantile estimation, Risk factors, Enterprise Risk Management, Accelerated algorithm, Nested Simulations.

INTRODUCTION

We denote by $X = (X_1, \dots, X_k)$ a random vector that can for example correspond to risk factors and consider a random variable Y which is a function of the risk factors : $Y = f(X_1, \dots, X_k)$. Our goal is to provide a way of computing the Value at Risk (VaR) of the random variable Y by simulations. Recall that for $\alpha \in]0, 1[$, $\text{VaR}_\alpha(Y) = q$ is the α -Value at Risk of Y if and only if it corresponds to the quantile

$$q = \inf\{x \in \mathbb{R}, P(Y \leq x) \geq \alpha\}.$$

We are interested in the general problem of estimating remote, but not extreme quantiles of Y thanks to accelerated simulations, in the case (frequent in risk management) when the function f is complex and very costly to evaluate, but when we have qualitative information on f . Charpentier and Oulidi (2009) explain the pros and the cons of the use of VaR optimization, in particular in terms of risk estimation. Our approach differs from the ones that one would use without information on f (please see conclusion

for a discussion on potential improvement of our method with tools that are useful in that case). For the insurance applications that we have in mind and that constitute our main motivation, $\alpha = 0.5\%$ is the regulatory quantile threshold and the variable Y corresponds to the net asset value of the insurance company after one year.

Since in applications, the computation of f requires some heavy computations, we are looking for *primary risk clouds* R such that the conditional Value at Risk equals the α -VaR :

$$(0.1) \quad \text{VaR}_\alpha(Y) = \text{VaR}_{\alpha'}(Y|(X_1, \dots, X_k) \notin R)$$

with $\alpha' = \frac{\alpha}{\mathbb{P}((X_1, \dots, X_k) \notin R)}$, and for $\beta \in]0, 1[$,

$$\text{VaR}_\beta(Y|(X_1, \dots, X_k) \notin R) = \inf\{x \in \mathbb{R}, P(Y \leq x|(X_1, \dots, X_k) \notin R) \geq \beta\}.$$

At a first stage, we shall assume that the function f is concave and continuous, and consider compact sets R . We shall construct nested sets R_u (see Sections 2 and 4) and we are looking for $u \in]0, 1[$ such that (0.1) is satisfied for $R = R_u$ but not for $R = R_{u'}$ with $u' > u$.

We are faced with two problems. First of all, we shall provide a practical criterion to recognize that a set R provides an estimation of $\text{VaR}_\alpha(Y)$ with a reasonable confidence. Secondly, we shall propose some ways of constructing nested sets R_u . To address these two questions of general interest, we try to use a non-specific vocabulary in Sections 1 and 2 and apply our methods to the insurance and risk management framework only from Section 3 on. Nevertheless, we now briefly expose the main motivation of our paper and the way it is positioned in the insurance and risk management literature.

The definition of economic capital in the new European insurance regulation Solvency II relies on the notion of 1-year equity distribution. Given the complexity of asset-liability dependencies bred by participating bonuses, dynamic surrenders and asset-liability management actions, it is not an easy task to obtain this distribution in closed form for life insurance portfolios. Primary simulations are carried out and correspond to the real-world evolution of the different risk drivers. After each primary simulation, one must derive the economic balance sheet of the company at time 1 year given what has occurred. To do this, secondary, (theoretically) market-consistent simulations are carried out, in order to compute a risk-neutral expectation. To estimate a 0.5% quantile of the net value of the company, this method is time consuming and difficult to implement in practice. From an actuarial point of view, it is clear that it is impossible to accomplish this task anyway, because of valuation issues in adverse scenarios, model risk and because of computation time limitations. We do not try to improve valuation techniques or to address model risk. This needs to be done but is beyond the scope of this paper. We take the modeling approach of insurance companies as they are, and we address the computation time limitation only.

A first approach could consist in using bagplots (see Rousseeuw, Ruts and Tukey (1999)). However, as mentioned in McNeil (2010), this approach performs well when risk factors follow an elliptic distribution and when the net asset value is linear in the risk factors. In that case, one could directly conclude that the Solvency Capital Requirement in the internal model would be the same as the one in the standard formula.

Besides, the linearity of the net asset value function is not verified and nonlinearities are important in practice (see Devineau and Loisel (2009b)). Outside of the elliptic assumption, bagplots can be constructed from numerical algorithms in dimension 2, but it seems not really feasible in higher dimensions, which is a problem for this kind of applications.

Devineau and Loisel (2009a) proposed a method to accelerate this algorithm and reduce computation times. In this early work, risk factors were assumed to follow a multivariate Gaussian distribution, and the convergence of the accelerating algorithm was only empirically tested.

In this paper, we first aim at proving that, under some conditions on the net value function that appear to be reasonable in practice, the accelerated estimator is consistent. Second, risk factors may have a non-Gaussian distribution (because marginal distributions may be heavy-tailed, or because the copula of risk drivers may exhibit strong tail dependence). They may also be simply obtained by simulation thanks to some only partially known commercial economic scenario generator. Therefore, we relax here the Gaussian hypothesis and allow the risk factors to have any copula and any marginal distributions. We explain how to implement our accelerating algorithm in both cases where the distribution of the risk factors is known or unknown. We expose and compare methods involving either geometric quantiles or copula density level curves or hyper-surfaces with so-called inverse methods. We also provide a verification algorithm that enables one to check that the algorithm is convergent, under the hypotheses made on the shape of the net asset value function. This verification possibility is important to reduce operational risk and model risk, and to build Solvency II internal models that meet the auditing standards of the second pillar of Solvency II (dedicated to internal control and audit of the processes). The efficient implementation of the verification algorithm leads to a nontrivial geometry problem, that we partially address in this paper. More precisely, to avoid too many computations, one must find a convex polytop with few extremal points and good enveloping or excluding properties with respect to the empirical risk driver scatter plot or with respect to a relevant level curve or hypersurface. Numerical illustrations on real-world saving portfolios are given, with two and three main risk factors (equity, interest rate, and property). Finally, we mention some ideas about dimension reduction when the number of risk factors is too high (which is the case in Solvency II if one keeps all QIS5 (Fifth Quantitative Impact Study) risk drivers) and about the way to tackle estimation risk due to Monte Carlo simulation errors. These two issues are briefly highlighted and their in-depth treatment is left for further research.

Our paper is organized as follows: in Section 1, we propose a way to estimate the Value at Risk with few calls to the function f , we prove that it is consistent and provide confidence intervals, then we recall the accelerating algorithm proposed in Devineau and Loisel (2009a) and show how to use the estimation to prove that the algorithm converges to the desired value. In Section 2, we explain how to get the frontiers of the subset in which secondary simulations can be avoided, both in the case where the distribution of the risk factors is known and in the case where it is not. In Section 4, we study efficient and pragmatic ways to construct convex polytops that are useful for the

verification procedure. Numerical illustrations are given in Section 3 in dimensions 2 and 3. The conclusion highlights potential future research developments.

1. AN ALGORITHM BASED ON PRIMARY RISK CLOUDS.

Recall that $X = (X_1, \dots, X_k)$ denotes a random vector of risk factors, that the risk indicator Y is a function of the risk factors : $Y = f(X_1, \dots, X_k)$ and that our goal is to provide a way of computing the Value at Risk (VaR) of the random variable Y by simulations.

1.1. Estimation of $\text{VaR}_\alpha(Y)$. For a given convex set R , given a n -sample $E_n = (X^1, \dots, X^n)$ of the random vector X , we estimate $\text{VaR}_\alpha(Y)$ by

$$\widehat{q}_R = C^{(j_\alpha)} \text{ where}$$

- (1) m is the number of elements of E_n that lie outside R ,
- (2) $\{Z^1, \dots, Z^m\}$ is the set of elements of E_n that lie outside R ,
- (3) for $j = 1, \dots, m$, $C^j = f(Z^j)$,
- (4) $C^{(1)} \leq C^{(2)} \leq \dots \leq C^{(m)}$ is the ordered sample of $\{C^1, \dots, C^m\}$
- (5) j_α is the smallest index j such that $\frac{j}{m} \geq \alpha'$ with

$$\alpha' = \alpha \times \frac{n}{m},$$

so that j_α is the smallest index j such that $j \geq \alpha n$.

We denote by \widehat{q} the standard estimator of $\text{VaR}_\alpha(Y)$ on a n -sample :

$$\widehat{q} = W^{(j_\alpha)}$$

with $W^{(1)} \leq \dots \leq W^{(n)}$ the ordered sample of $W^1 = f(X^1), \dots, W^n = f(X^n)$.

As an example, if $n = 5000$, $\alpha = 0.5\%$, $m = 500$ then $\alpha' = 5\%$ and in order to compute \widehat{q}_R , only 500 calls to the f function are needed (rather than 5000 if one considers the entire sample E_n).

The question is: how to choose R so that $\widehat{q}_R = \widehat{q}$? The idea is that points in E_n that contribute to \widehat{q} are not in the "center" of the cloud. This will be made precise in Section 2.

1.2. Conditions of convergence. First of all, if f is a concave function then it may be checked with few calls to the function f that $\widehat{q} = \widehat{q}_R$.

Theorem 1.1. *Let R be a given convex set. Denote $M_R = \inf_{x \in R} f(x)$. With the above notations, let ℓ be the number of elements of $\{C^1, \dots, C^m\}$ such that $C^j < M_R$. If $\ell > \alpha n$ then $\widehat{q} = \widehat{q}_R$ and thus \widehat{q}_R is a consistent estimator of $q = \text{VaR}_\alpha(Y)$.*

Moreover, if f is a concave function, the condition $C^j < M_R$ may be checked with a finite number of f computations.

Furthermore, fix a confidence threshold $\beta \in]0, 1[$ (namely, $\beta = 5\%$), and let z_β be the $1 - \beta$ quantile of the standard Normal distribution and \widetilde{j} be the smallest index j such that

$$(1.2) \quad \frac{j}{m} \geq \alpha' - z_\beta \sqrt{\frac{\alpha'(1-\alpha)}{m}}$$

(recall that $\widehat{\alpha}' = \alpha \frac{n}{m}$) then

$$\mathbb{P}(C^{(\bar{j})} \leq q) \xrightarrow{n \rightarrow \infty} 1 - \beta.$$

In other words, with an asymptotic $1 - \beta$ confidence, $C^{(\bar{j})}$ is a lower bound for $\text{VaR}_\alpha(Y)$.

Proof. It is well known (see for example Van der Vaart (1998)) that \widehat{q} is a consistent estimator of q (provided that q is not an extreme quantile, i.e. α is fixed with respect to n). Our condition that there are more than αn elements C^j of the sample outside R with the property that $C^j < M_R$ clearly implies that $\widehat{q} = \widehat{q}_R$.

For the second part, recall that a concave function on a convex set takes its minimum on the extremal points of the convex set. Thus, let S be a convex polytop, such that $R \subset S$ then S has finitely many extremal points on which the minimum of f on S is reached. Then, $M_S := \inf_{x \in S} f(x) \leq M_R$. In order to compute M_S , we have to evaluate f at the extremal points of S and it is sufficient to have more than αn indexes j such that $C^j \leq M_S \leq M_R$.

The last part follows from results on ordered statistics which give asymptotic confidence intervals for quantiles (see Van der Vaart (1998) for example). Remark that (1.2) is equivalent to

$$\frac{j}{n} \geq \alpha - z_\beta \sqrt{\frac{\alpha(1-\alpha)}{n}}.$$

□

The concavity condition on f may be weakened in various directions. For example, in dimension 2, if the function f is concave on the upper half space and coordinate-wise decreasing on the other half plane, Theorem 1.1 remains valid. We experienced a real-world case in which this situation occurred. Besides, in some special cases where complex, partial hedging programs are used, it may be possible that the concavity assumption is only valid on certain sub-domains.

Of course, once Theorem 1.1 has been proved, the main question is to find an adequate convex polytop. We shall propose some ways to obtain it in Sections 2 and 4. Beforehand, we now show in Subsection 1.3 how Theorem 1.1 may be used to validate the so-called accelerated algorithm described in Devineau and Loisel (2009).

1.3. An empirical algorithm. In Devineau and Loisel (2009) an algorithm has been proposed to estimate q with a n -constraint (minimize the number of calls to f) in an insurance context. We shall briefly describe it and show how Theorem 1.1 may be used to validate it.

Recall that in higher dimension the notion of outer points (or outers) may be defined in several ways (in dimension 1, one may consider as $x\%$ outers smaller or greater values of the sample). In Section 2, we describe how geometric quantiles and depth functions may be used to define outers. In Devineau and Loisel (2009) the level curves of the density levels (of a Gaussian vector) were used for this purpose.

According to a definition of outers, the algorithm proceeds as follows. x is fixed (for $\alpha = 0.5\%$, $x = 2$ seems to be a convenient choice from an empirical point of view).

- Let R_1 be a subset of E_n such that the $x\%$ outers are outside of R_1 .
- Compute f on these $x\%$ outers. A first estimation of q is \widehat{q}_{R_1} .
- Consider a smaller subset $R_2 \subset R_1$ such that the $2x\%$ outers are outside R_2 .

- Compute \widehat{q}_{R_2} . If $\widehat{q}_{R_1} = \widehat{q}_{R_2}$ then stop; else continue with the 3x% outliers that are outside R_3 ...
- The algorithm stops as soon as $\widehat{q}_{R_k} = \widehat{q}_{R_{k+1}}$ for some k or at the end if we have used the whole sample without stopping.

We denote by R_k the stopping cloud ($R_k = E_n$ if we have used the whole sample). In case $R_k \neq E_n$ (which is of course the interesting case), Theorem 1.1 may be used to verify that R_k provides a consistent estimator of the Value at Risk, if we assume that f is concave or has some properties of local concavity. Indeed, consider a convex polytop $S \supset R_k$. Compute (with as many calls to f as the number of extremal points of S) $M_S = \inf_{x \in S} f(x)$.

If $\widehat{q}_{R_k} \leq M_S$ then we know that $\widehat{q}_{R_k} = \widehat{q}$.

In Section 5, we compute the probability that the algorithm stops while it should not.

2. DIFFERENT WAYS TO DETERMINE EXTREMAL PRIMARY POINTS

We present here different ways to define contours to exclude points that are inside this border from secondary computations. A first method corresponds to using contour plots of the joint density of the primary risk driver vector. A clear drawback is that it has to be estimated if we do not know the distribution of these primary risk drivers (for example if they are obtained from a black-box proprietary software). Besides, this may lead to non-convex domains, for example in the case of Gaussian risk driver vectors with non-trivial Clayton copulas. Another possibility is to use geometric quantiles (see for example Chaouch and al. (2008)), or a so-called inverse method inspired from geometric quantiles. These methods lead to more or less elliptical contours or ellipsoids in higher dimensions, even if the density contours are not elliptic. Another possibility would be to use depth functions. However this concept works well only in dimension 2, and in higher dimensions for linear functions of elliptic vectors. If one wants to better take into account the shape of the risk cloud, one could modify the depth function in order to better stick to the periphery of the risk cloud. Nevertheless, these methods are not so useful for our problem for two main raisons: the lack of feasibility in higher dimensions, and the non-convexity of the domains, which prevents us from using our verification algorithm. So as a conclusion, inverse methods seem to be a good compromise that gives us the opportunity to check our result when the net asset value has desirable properties.

2.1. Density contours. In order to accelerate Nested Simulations calculation, it is necessary to measure the risk associated to each simulation and to order them. Thus, most adverse simulations in terms of solvency are considered. This classification of scenarios leads to build an execution area with the points that seem the most risky. Remember that in the case of an elliptic distribution of risk factors (Gaussian distribution for example), it is possible to build this area by plotting contours of the simulations cloud (each point of the cloud being defined by the value of the risk factors for each simulation). Scenarios to be studied first are the ones inside the execution area. These scenarios lead to the values of equity which are located in the distribution tail, such as the quantile that leads to the economic capital is one of these scenarios (see Devineau and Loisel (2009)).

The key idea is to select scenarios that lead to small equity values at time $t = 1$ just

with the information given by risk factors. If the risk factors distribution is known, it becomes easy to define the execution area. Indeed, it is possible to calculate density contours directly from the density function of the distribution, in such a way that every point with a density lower than a given threshold is a part of the execution area. Therefore this area is defined by the points on the border of the cloud, and the threshold can be determined such as the area is made by a given number of points. In the case where the distribution is unknown, it is still possible to use a norm to define the contour of the cloud and focus the analysis on the most adverse simulations.

If the distribution of risk factors is unknown as it is often the case (if risk factors are deduced from an existing proprietary scenario table, the copula of risk factors may be unknown), the density contour of the cloud can be determined with an empirical approach. This leads to a first limit of the technique that consists in building an execution area thanks to the density contours. Besides, we shall see further that if this process is efficient for an elliptic distribution, it is no longer the case for non-elliptic distributions (such as distribution with a Clayton copula as we study later in this section).

To generalize the construction of the execution region in a "distribution-free" framework, we have considered other alternatives: the calculation of a multi-dimensional quantile first, then the introduction of a depth function.

2.2. Geometric quantiles. To characterize the contours of a given distribution, it may be interesting to estimate the quantiles of the distribution and to consider the lines of constant density. However, the generalization of the concept of quantile in the multi-dimensional framework is not straightforward because of the lack of total ordering criterion in dimension strictly higher than 1. The article by Barnett (1976) proposes some techniques for ordering multivariate observations and is considered as a true starting point for the study of bivariate quantiles. Recently, two main approaches have been developed to define multivariate quantiles that are invariant under rotation and / or affine transformation. On the one hand, the geometric quantiles, introduced by Chaudhuri (1992), are defined as solutions of a problem of a loss function minimization, by generalizing a known result in dimension 1. On the other hand, another approach is based on the depth function.

Recall that in dimension 1, the quantile of a random variable X is defined using the inverse of the distribution function of this variable. If we denote this function F , for $u \in [0, 1]$, the u -quantile is defined as:

$$Q_F(u) = F^{-1}(u) = \inf \{y/F(y) \geq u\}$$

It may be proved that $Q(u)$ is the solution of a minimization problem (see for example Chaouch et al. (2008)):

$$(2.3) \quad Q(u) = \operatorname{argmin}_{x \in \mathbb{R}} \mathbb{E} \{ \phi(u, X - x) \},$$

with $\forall (u, t) \in]-1; 1[\times \mathbb{R}$, $\phi(u, t) = |t| + ut$,
so that $Q(u)$ is also the root of the equation

$$(2.4) \quad \mathbb{E}(\operatorname{Sign}(x - X)) - u = 0.$$

The definition of quantile may be generalized in higher dimension in the following way. Consider a random vector X which takes its values in $\mathbb{R}^d, d \geq 2$. Then let

$$\forall(u, t) \in B^d = \{u \in \mathbb{R}^d / \|u\| \leq 1\} \times \mathbb{R}^d : \phi(u, t) = \|t\| + \langle u, t \rangle$$

and the u -quantile is defined as follows ($\|u\| \leq 1$):

$$\begin{aligned} Q(u) &= \operatorname{argmin}_{x \in \mathbb{R}^d} \mathbb{E} \{ \phi(u, X - x) \} \\ (2.5) \quad &= \operatorname{argmin}_{x \in \mathbb{R}^d} \mathbb{E} \{ \phi(u, X - x) - \phi(u, X) \}. \end{aligned}$$

Subtracting $\phi(u, X)$ in the equation does not change the solution of the minimization problem but ensures the existence of the expected value regardless of the expectation of X (that may be infinite). The quantile is the solution of the following equation that generalizes (2.4):

$$(2.6) \quad \mathbb{E} \left(\frac{x - X}{\|x - X\|} \right) - u = 0.$$

In practice, we construct an empirical estimator of the quantile value with n observations (X_1, X_2, \dots, X_n) of the random vector X that constitute the cloud. This estimator $Q_n(u)$ is defined as a solution of the following equation:

$$(2.7) \quad \frac{1}{n} \sum_{i=1}^n \left(\frac{x - X_i}{\|x - X_i\|} \right) - u = 0.$$

The vector u provides information on the position of the quantile: if the norm of u is close to 1, the quantile is "extreme", that is to say away from the median. Furthermore, the direction of the vector u gives us an indication of the position of the quantile in regard of the median.

We will look at first two distributions of risk factors in dimension 2 to implement and understand the technique and analyze its limitations. Then we shall see in Section 3 a case study in dimension 3 by considering three risk factors and the allocation of economic capital on a saving product.

2.3. Practical use of the geometric quantiles. The geometric quantiles we have just defined are not invariant through an affine transformation. So that, if the various components of the vector X do not have comparable orders of magnitude or if they do not show similar variations, the quantile results are not satisfactory. If we assume that risk factors are distributed with a Gaussian copula, it is trivial to standardize risk factors if necessary to circumvent the problem. If this is not the case, we must use another estimation technique, called transformation-retransformation technique (see Chaudhuri and Sengupta (1993)). This technique in particular makes the quantile invariant under affine transformation.

It proceeds as follows. Consider a vector X of n observations in \mathbb{R}^d with $n > d + 1$ and $\alpha = (i_0, i_1, \dots, i_d) \subset \{1, \dots, n\}$. The transformation matrix of order $(d \times d)$ can be defined as:

$$X(\alpha) = (X_{i_1} - X_{i_0}, \dots, X_{i_d} - X_{i_0})$$

This matrix is used to transform all the points X_j such that $j \notin \alpha$. The vector α is determined so that the matrix $[X(\alpha)^T] \Sigma^{-1} X(\alpha)$ is close to a diagonal matrix, Σ being

the variance-covariance matrix of the random variable X , and M^T the transpose of the matrix M . This constraint is equivalent to choose α in order to minimize the ratio of the quotient between the arithmetic and the geometric mean of the eigenvalues of the matrix $[X(\alpha)^T] \widehat{\Sigma}^{-1} X(\alpha)$, where $\widehat{\Sigma}$ is an empirical estimator of Σ . This constraint is also equivalent to minimize the ratio between the trace and the determinant of this matrix. The transformation is done through the following formula, $\forall j \notin \alpha$:

$$X_j^\alpha = [X(\alpha)^T] X_j$$

The dimension of the transformed vector X^α is $n - (d + 1)$. It is also necessary to transform the vector u while ensuring it remains in the unit sphere. The vector $v(\alpha)$ is defined as the transform of the vector u :

$$v(\alpha) = \begin{cases} \frac{\|u\|}{\|X(\alpha)^{-1}u\|} X(\alpha)^{-1}u & u \neq 0 \\ 0 & u = 0 \end{cases}.$$

The estimator of the geometric quantile with order $v(\alpha)$ is then calculated from observations X_j^α such that $j \notin \alpha$ using the equation 2.7. The estimate of this quantile is denoted by $R_n^\alpha(v)$. The retransformation step gives the final estimator of the quantile. This estimator of order u is given by $Q_n^\alpha(u) = X(\alpha)R_n^\alpha(v)$. The reader can turn to Chaouch, Gannoun and Saracco (2008) for theoretical validation of this methodology.

To implement this technique on a given cloud of points, a possible method is to use a Newton-Raphson algorithm. We start by creating a series of N vectors $(u_k)_{1 \leq k \leq N}$ in B^d . Each vector u can be associated with a quantile value calculated using a step by step solving algorithm: the principle is to compute approximations of quantiles by recurrence and stop the algorithm when two successive approximations are close enough. Assuming the initial approximation is composed of marginal means of the N vectors u , the approximation of order $m + 1$ is obtained from the approximation of order m in the following way:

$$Q_n^{m+1}(u) = Q_n^m(u) + \Phi^{-1}\Delta,$$

$$\text{with } \Phi = \sum_{i=1}^n \frac{1}{\|X_i - Q_n^m(u)\|} \left(I_2 - \frac{X_i - Q_n^m(u)}{\|X_i - Q_n^m(u)\|} \left(\frac{X_i - Q_n^m(u)}{\|X_i - Q_n^m(u)\|} \right)' \right)$$

$$\text{and } \Delta = \sum_{i=1}^n \frac{X_i - Q_n^m(u)}{\|X_i - Q_n^m(u)\|}.$$

By choosing N sufficiently large and a consistent norm u for the N vectors, it is possible to construct a quantile curve for a given distribution.

Let us implement this methodology on the simple example of a distribution of n simulated pairs in \mathbb{R}^2 , drawn from the law (X_1, X_2) . X_1 and X_2 are correlated and we assume that the two series have different orders of magnitude. The transformation-retransformation algorithm allows us to plot the outline of the cloud of points by calculating the quantiles for different directions of the vector u . It appears in the graph below that the quantile curve takes the shape of the distribution of pairs, and identifies the worst scenarios (those that are not included in the curve).




FIGURE 1. quantile curve obtained using the algorithm of transformation-retransformation on an elliptical cloud

2.4. Inverse method. The major advantage of the geometric quantile approach lies in the fact that it is a “distribution-free” method. Indeed, the norm used to classify scenarios works under the hypothesis of the normality of the risk factors copula. But this assumption is not always checked. Risk factors retro-calibrated thanks to “standard” diffusion models will not necessarily present a Gaussian copula.

However, the Transformation-Retransformation (TR) approach used above is very costly in calculation time in the way that it solves an equation step by step with a Newton-Raphson algorithm. Besides, there is no need from a strictly operational point of view in insurance to build the border of the execution area to determine which scenarios are inside, the method can be used just to classify scenarios and to study the extreme ones. Indeed, if we reverse the problem, we just have to associate a threshold vector u for each simulated pair. Assuming that the executing area satisfies good properties (including the fact that the norm of the quantile is a monotonic function of the vector threshold), it is possible to classify the scenarios according to their associated threshold, and thus characterize the notion of periphery of the cloud.

The first step is to determine, in the same way as in the TR algorithm, a transformation matrix $X(\alpha)$ with dimension $d \times d$. Let us consider every simulated pair as a value of a transformed quantile. We can deduce the empirical estimator of R^α :

$$R_n^\alpha = X^{-1}(\alpha)X.$$

From these transformed points that we consider as values of quantiles, denoting X^α the vector of transformed random variables, we deduce for each point of the distribution

a vector v such that:

$$v = \frac{1}{n} \sum_{i=1}^n \frac{R_n^\alpha - X_i^\alpha}{\|R_n^\alpha - X_i^\alpha\|}.$$

The norm of the vector u is the same as the vector v by definition. The vector v we have just defined can be calculated in order to give a value of the degree of risk associated to each scenario.

In this way, with the same distribution as introduced above, the scenarios with an associated threshold of high norm are correctly located on the outskirts of the cloud, as shown in the graph below.




FIGURE 2. The points with threshold vectors of higher norm are pointed in red

2.5. Clayton copula. The density contours given by threshold vectors are often elliptical. We present the case of vectors with Gaussian marginal distributions linked with a Clayton copula. The implementation of the transformation-retransformation algorithm leads to an elliptical execution area border as shown on figure 4. The density contour does not fit the shape of the cloud whose points in the lower left quadrant are clustered around the diagonal, while those in the upper right quadrant are scattered around the same diagonal. We thus observe an operational limit of geometric quantiles and can deduce that it will be difficult to plot the contours of non-convex clouds (such as the cloud formed by the Clayton copula) with this method.

The drawing of both contour lines obtained from the copula density and quantile lines obtained by the method of transformation-retransformation clearly shows the error due to the non-convexity of the cloud. The poor quality of adjustment induced by a Gaussian copula hypothesis results from a failure of ellipticity of a Clayton copula




FIGURE 3. Geometric quantile contour in black versus density contour in red for a Clayton copula

cloud.

Using the inverse method on the same distribution, the same result appears.

Prioritization of simulations with the norm of the threshold vectors in the geometric quantile approach does not provide satisfactory contours of the cloud, the execution area border being convex.

2.6. Depth function. To answer to the issues mentioned above and find a totally "distribution-free" method, it can also be useful to consider another approach of multivariate quantiles. Along with the geometric quantiles, the notion of depth function has been developed in recent years to characterize the quantiles of a distribution in multi-dimensions.

The depth of a point that belongs to a cloud is used to quantify the fact that this point is on the outskirts of the cloud or whether it is inside of it. The principle is to determine if such a point is surrounded by other points or not.

To quantify the depth of each point, we introduce the half-space defined by a point y of the cloud and a vector u of dimension d . This area is characterized by the following equation:

$$H_{y,u} = \{x \in \mathbb{R}^d / u'x \geq u'y\}$$

The depth function then associates to each point of the plane the minimal probability to find other points in the half-space defined above. If we consider the empirical probability calculated from the points of the cloud (x_i), the depth function is defined as




FIGURE 4. Determination of the 100 points with quantiles of highest norm for the Clayton copula

follows:

$$p(y) = \inf_{u \in \mathbb{R}^d} \{P(H_{y,u})\}$$

$$\widehat{p}(y) = \min_{u \in \mathbb{R}^d} \frac{1}{N} \sum_{i=1}^N 1_{\{u'x_i \geq u'y\}}.$$

It is then possible to classify the points of the cloud by increasing depth and thus characterize the extreme scenarios. If this method is effective for convex distributions, it is unlikely that it works on non-convex distributions, since the execution area built with this method is an intersection of half-spaces. The region border is therefore convex by definition. The prioritization of scenarios by increasing depth can not solve the issues raised by the use of geometric quantiles. Indeed, the non-convexity implies the existence of points located on the edge of the cloud and yet with a depth that can be high, which is non coherent with the definition we want to give to the depth function. This phenomenon is explained by the graph below, showing the example of the Clayton copula.

On the graph appear on one side the line that defines a half-plane $H_{y,u}$ for a point y located in a hollow of the cloud and for a vector u such that the minimum of the empirical probability of this half-space is reached, and on the other side a circle around some of the points that belong to this half-space. This situation highlights the fact that by construction, the depth of the studied point does not match its actual position relative to the cloud. Indeed we observe a point on the periphery of the cloud whose depth is high. Moreover, we note that the points of lowest depth pointed in red on the




FIGURE 5. Calculation of the depth of a point on the outskirts of Clayton copula

graph are all located in convex areas of the cloud.

In the case of a non-convex cloud, the conclusion is that such an approach systematically fails. Besides, this approach is difficult to implement in dimension larger than 3 because of the complexity of estimating the depth for each point.

2.7. Modified Depth Function. We present in this section a new measure: the modified depth function. We are changing the definition of depth so that a point in a hollow is not penalized as it is the case in the previous example with the standard depth definition. We calculate the probability to find other points in a parabola instead of considering a half-space. Noting the modified depth p^* and $H_{y,u}^*$ the parabola that has its origin in y and with u as its directional vector, we define this function as follows:

$$p^*(y) = \inf_{u \in \mathbb{R}^d} \{P(H_{y,u}^*)\}$$

$$\widehat{p}^*(y) = \min_{u \in \mathbb{R}^d} \frac{1}{N} \sum_{i=1}^N \mathbf{1}_{\{u'(x_i - y) \geq \|(x_i - y) \cdot u^\perp\|^2\}}$$

The use of this indicator very close to the depth previously introduced can greatly improve the quality of adjustment in some non-convex cases like in the case of the Clayton copula.

On Figure 2.7, the parabola is the curve that defines the modified depth. The items that appear in the circle and were taken into account in the calculation of the standard depth are not part of the parabola, which implies that the modified depth of points in the non-convex areas is sufficiently low to consider these points as outliers (in the periphery of the cloud). Pointing the points of lowest modified depth in red, the outline of the

cloud appears.




FIGURE 6. Determination of the extreme points of the Clayton copula with modified depth

It is possible to extend the parabola by adding a multiplying term in the constraint of the indicator appearing in the calculation of the modified depth. A deeper theoretical analysis on the properties of the modified depth function is necessary to study its domain of applicability, as well as its performance for elliptic and non-elliptic risk drivers. This will be the purpose of a further work.

3. NUMERICAL APPLICATION

We detail here the application of the methods presented in the previous section. Studying the case of a "standard" insurance company, the goal of the following results is to compare the use of a basic accelerator (using density level curves) and the use of geometric quantiles (or depth function).

3.1. Description of liability and models. The portfolio that we consider in this study is a saving portfolio with low guaranteed interest rate. We have projected this portfolio using an internal model that performs ALM stochastic projections and the calculation of equity after one year. This projection tool enables us to model the profit sharing mechanism, as well as dynamic lapses of the policyholders when the interest rates handed out by the company are deemed insufficient in relation to the reference interest rate offered by the competitors.

At a first stage, we consider the stock and interest rate as related risks at a first stage (2-dimensional study). At a second stage, stock implied volatility risk is introduced,

which leads to a 3-dimensional study. The tables of economic scenarios that are used were updated on December 31th, 2009.

The net asset value at time 1 is a function of the considered risks. For real portfolios, it has been observed that a concavity or a partial concavity hypothesis for f is reasonable (recall that the standard formula uses an affine function).

3.2. Results for 2 risks. First, let us consider only the stock and the interest rate risks. We consider in this section two different acceleration algorithms. The first one uses the density level curves to define outliers whereas the second one uses geometric quantiles. In this example, the risk vector is assumed to be Gaussian. Let us note that many of the models used by the insurance business are Gaussian (are Gaussian based). Nevertheless, these Gaussian hypotheses must be relaxed in order to take into account non Gaussian dependencies between risks as well as heavy tailed distributions that are relevant for insurance business. This is why we propose to use distribution free approaches such as geometric quantiles or depth functions.

In the table below we compare the execution of the accelerated algorithm (as described in Section 1.3) when the outliers are defined by using density level curves with respect to geometric quantiles. For these executions, $n = 5000$, $\alpha = 0.5\%$ and $x = 2$. We give the number K_{ma} of points that lie outside the stopping cloud R_k and the number of calls to function f before the algorithm stops.

Method	Density level curve	Geometric quantile
K_{ma}	153	150
Number of calls to f	300	300

TABLE 1. Convergence of the accelerator algorithm in dimension 2

Both methods lead to equivalent results. Indeed, geometric quantiles capture well the shape of the cloud for elliptical distributions.

The table below illustrates the confidence bounds for the VaR as described in Section 1, for $n = 5\,000$, $n = 50\,000$, $n = 500\,000$. Of course, and particularly if the size n

n	5 000	50 000	500 000
Estimated VaR (thousand €)	-39 375	-41 787	-43 058
rank	25	250	2 500
Lower 5% confidence bound (thousand €)	-44 722	-43 840	-43 173
rank	20	236	2 458

TABLE 2. Confidence bounds

of the sample is not too large, the use of a confidence bound rather than a pointwise estimation should be preferred.

3.3. Results for 3 risks. In this section we introduce a new risk: the stock implied volatility risk.

Due to the introduction of the stock implied volatility risk, the algorithm converges more slowly using geometric quantiles rather than the level curves. We conclude from




FIGURE 7. Distribution of risk factors - red points represent highest geometric quantiles and blue ones are points on highest density level curves

Method	Density level curve	Geometric quantile
K_{ma}	105	292
Number of calls to f	300	400

TABLE 3. Convergence of the accelerator algorithm in dimension 3

this observation that in the case where the assumptions on the distribution of risk factors do not raise an issue, geometric quantiles are not useful to locate extreme scenarios in terms of equity value. The use of density level curves is therefore highly efficient as well as easy to implement for an elliptical distribution of risk factors.

Besides, we note that the results obtained with the modified depth function on this product lead to similar results to those obtained with the geometric quantiles. We infer that these methods are useful when the distribution of risk factors no longer meets certain assumptions (such as the ellipticity). But they do not improve the results if these assumptions are met.

4. CONSTRUCTION OF THE CONVEX VERIFICATION SUBSET

In general, it might be quite difficult to construct a verification convex polytop or polygon (introduced in Subsection 1.2) that is not too far from the border of R_u (in order to avoid to include or exclude too many points in comparison to R_u) and does not have too many vertices (in order to limit secondary, time costly simulations to evaluate the position). Here we present a simple, pragmatic way to address the 2-dimensional case where the distribution of the risks vector is elliptical. It might be extended in

practice to geometric quantiles whose level curves are not far from being elliptical in that case. Recall that the accelerating algorithm starts from outers of the cloud and goes back towards its center. Consider the example studied in the previous section, with two risk drivers (asset and interest rate risks). The first step is to use dilatations to restrict without loss of generality R_{it} 's defined by circles. In that case, the algorithm converged after 3 rounds of 100 simulations: the result after considering 300 outlying simulations (keeping points that are outside the green circle in Figure 8 below) is the same as the one after 200 (out of 5000) outlying simulations (keeping points that are outside the blue circle). Therefore, it makes sense to use a convex verification polygon that lies in between the two circles. The radiuses of the two circles are very close to the respective theoretical 94% and 96%-quantiles of the radius of a standard Gaussian vector ($r_1 = 2.37$ and $r_2 = 2.54$). Therefore the minimal number of vertices for the regular polygon included in the outside circle with first extremal point $(r_2, 0)$ to contain the circle of radius r_1 is given by

$$\lceil \frac{\pi}{\text{Arccos}\left(\frac{r_1}{r_2}\right)} \rceil = 9$$

in our case. The convex polygon and the two circles are shown in Figure 8.




FIGURE 8. The two circles and a convex verification polygon in the model of Subsection

Note that we considered here Gaussian risk factors to illustrate our method because this corresponds to models used by insurance companies. Nevertheless we encourage the use of distributions with heavier tails for risk drivers that feature this kind of pattern. In that case, a few outers might be far away from the rest of the risk cloud, and the circles are likely to have larger radiuses and to be less close to each other. In that case one would need fewer vertices to build a verification convex polygon between the two circles. However, it might be interesting to increase a bit the number of vertices in order to get a verification convex polygon that better sticks to the outer circle. To

illustrate this, we present an example in which (standardized, i.e. decorrelated) risk drivers are drawn from an isotropic distribution where the radius of the 2-dimensional vector is distributed as a stable distribution with infinite mean (this might be the case for operational risk as mentioned by McNeil et al. (2005)). Let X be a centered 2-dimensional isotropic stable random vector with characteristic function $\exp(-\gamma_0^\alpha |u|^\alpha)$.

The amplitude (radius) of X is defined by $R = \sqrt{X_1^2 + X_2^2}$. From Nolan (2005), when the dimension is $d = 2$ and when $\alpha = 1$, the c.d.f. of R is given for $r > 0$ by

$$F_R(r) = 1 - \frac{\gamma_0}{\sqrt{\gamma_0^2 + r^2}}.$$

Stable distributions may be simulated thanks to John Nolan's STABLE program (available at <http://academic2.american.edu/~jpnolan/stable/stable.html>).

For $\gamma_0 = 0.15$, the 94% and 96%-quantiles of R are respectively around 2.5 and 3.75. In that case the minimal number of vertices for a regular polygon to lie in between the two circles falls down to 4, but one may choose to keep it equal to 9 for example in order to enlarge the verification zone. The risk cloud is shown in Figure 9. One outlier in the upper left-hand corner prevents us from seeing anything. We thus logarithmically rescale the radius of outliers (outside the circle of radius 3.75) in order to obtain a transformed graph that is easier to read. The two circles and two convex polygons are shown in Figure 10 (after rescaling the radius of outliers). Note that in that case it might be difficult to value the position of the company in the few very extreme scenarios, because the world would be very different in those cases, and evaluation would obey different rules.




FIGURE 9. The risk cloud in the stable isotropic case without rescaling

In a more general setting, once one of the above-mentioned techniques (depth function, inverse methods, geometric quantiles, density level curves) has enabled us to avoid to evaluate the position of the insurance company at a certain number of points, it is possible to quickly construct the convex hull of those points thanks to several




FIGURE 10. The two circles and a convex verification polygon in the stable isotropic case (after rescaling the radius of outers with a logarithmic transform)

convex hull algorithms. Note that of course, if depth function is used in the case of Clayton copula risk clouds as the one of Figure 3, the outers are not outside a convex subset, and consequently the verification algorithm will not be as efficient as in the case of convex R_u 's.

Here we show an example with 200 points randomly drawn on a sphere in order to get something comparable to our problem in dimension 3: recall that in our case with 5000 primary simulations, we have 200 points between two spheres that delimit the zones of non-execution after step $K \geq 1$ and step $K+1$. So our situation is not too far from the one where 200 points are randomly drawn on the outer sphere (because all the other points inside in inner sphere are likely to be included in the convex hull of the 200 points drawn on the outer sphere (or in between the two spheres)). The convex hull of the 200 points is determined thanks to the QuickHull algorithm applet developed by Tim Lambert and available at <http://www.cse.unsw.edu.au/~lambert/java/3d/hull.html>. The convex hull is shown in Figure 11. The number of vertices is still reasonable (a few hundreds) for the verification algorithm in that case.

5. PROBABILITY THAT THE ALGORITHM STOPS ALTHOUGH IT SHOULD NOT

Insurers cannot do more than 5000 simulations in many cases. The SdS accelerator can improve computation times if things go well (if we have enough information on the characteristics of the net asset value function). If it is not the case, then the algorithm may take time to converge, and after more than 10 or 20 iterations, even if the algorithm finally converges, the user should be cautious, and in that case it could be necessary to run the 5000 points if doubts are present. So in reality what do not want to occur too often corresponds to events during which the algorithm stops very quickly (after a few runs of 100 points, fewer than 5 or 10, say) despite the fact that the target empirical quantile has not been reached yet. We aim at showing that, even in the case where we




FIGURE 11. The convex hull of 200 points randomly drawn on the sphere in dimension 3 (obtained from Tim Lambert's Java Applet)

have absolutely no idea of what the net asset value function looks like, the probability of such bad events is quite small.

Let R_k be the rank (between 1 and 5000) of the 25^{th} smallest value in sub-sample number $k \geq 1$ that contains all the points of the first k rounds (i.e. $100k$ if we select 100 points per round). We assume that samplings are without replacement and independent (apart from their without replacement characteristics).

For $25 \leq r \leq 4925$, we have

$$P(R_1 = r) = \frac{C_{r-1}^{24} C_{5000-r}^{75}}{C_{5000}^{100}}.$$

Given that $R_1 = r \leq 4825$, we have that

$$P(R_2 = r | R_1 = r) = \frac{C_{5000-r-75}^{100}}{C_{5000-100}^{100}}.$$

From the total probability formula, the probability that the accelerating algorithm inappropriately stops after 2 iterations is

$$q_2 = P(R_2 = R_1 \geq 26) = \sum_{r=26}^{4825} \frac{C_{r-1}^{24} C_{5000-r}^{75}}{C_{5000}^{100}} \frac{C_{5000-r-75}^{100}}{C_{4900}^{100}}.$$

With Maple, one gets that $q_2 = 5.363 \times 10^{-9}$: it would be bad luck to stop inappropriately after the second run!

After 9 runs, for $25 \leq r \leq 4125$, we have

$$P(R_9 = r) = \frac{C_{r-1}^{24} C_{5000-r}^{875}}{C_{5000}^{900}}.$$

Given that $R_9 = r \leq 4025$, we have that

$$P(R_{10} = r | R_9 = r) = \frac{C_{5000-r-875}^{100}}{C_{5000-900}^{100}}.$$

From the total probability formula, the probability that $R_{10} = R_9 \geq 26$ is

$$\tilde{q}_{10} = P(R_{10} = R_9 \geq 26) = \sum_{r=26}^{4025} \frac{C_{r-1}^{24} C_{5000-r}^{875} C_{5000-r-875}^{100}}{C_{5000}^{900} C_{5000-900}^{100}}.$$

With Maple, one gets that $\tilde{q}_{10} = 6.94\%$: this probability is too high: it means that if we are not sure about the hypotheses and if the algorithm converges after 10 runs, it might be useful to have a look at the estimations of VaR_α for $99\% \leq \alpha \leq 99.5\%$, say. The probability to choose the empirical VaR_β with $\beta < 99\%$ (instead of $VaR_{99.5\%}$) is $P(R_{10} = R_9 \geq 51)$ and is equal to 0.448%. Similarly, the probability to choose the empirical VaR_β with $\beta < 99.25\%$ (instead of $VaR_{99.5\%}$) is $P(R_{10} = R_9 \geq 38)$ and is equal to 1.88%. Note that the principal terms that contribute to this value are the ones for $100 \leq r \leq 300$.

Let us now consider an intermediate case that occurs quite often: convergence after 5 runs. The probability that the algorithm gives the same 25 worst scenarii in the fourth and fifth runs is $P(R_5 = R_4 \geq 26)$ and is equal to 0.323%, which is acceptable. Note that this is higher than the probability that the algorithm inappropriately stops exactly at the fifth run, because the algorithm could have stopped earlier.

Note that in reality, when we know something about the net asset value function, R_1 is likely to be small (much smaller than in the case where we do not anything about the function, because we have better chances to choose bad scenarios with some information than without any). Without information, as $P(R_2 = r | R_1 = r)$ is decreasing in $r \in [25, 4825]$, the conditional probability to have $R_{k+1} = R_k = r \geq 26$ given that $R_k = r$ decreases with r : this is good news, because the probability to have a problem is larger when the error is quite small. Besides, in the case with information, the probability to have smaller values of R_k is higher than in the case without information. One could think that the probability of failure of the algorithm increases in the information we have about the net asset value function, which would be very bad. Actually, with some information, one has much better chances to find new small values at the k -th run than in the case without information. To take this into account, one could use the fact that one can use the values of the first runs and the characteristics of the net asset value function to choose the next points and to bound the probability of failure of the algorithm (using for example Lipschitz parameter of the net asset value function).

The probabilities computed here are dimension-free, for 5000 primary scenarios. However, it is clear that if the dimension is high, the number of points needed to get accurate enough results becomes very large.

6. CONCLUSION AND FURTHER RESEARCH

We have provided a way to accelerate the nested simulation method in Solvency II and more generally to accelerate the estimation of quantiles of complicated functions of random vectors with some patterns, as well as a verification process. The results were mainly presented on real-world partial internal models of insurance companies, with typical models, parameters, evaluation techniques and number of simulations that are currently used by most European companies. Of course, we plea for the use of at least 1 million primary simulations to estimate the 99.5% VaR of the net asset value function

at time 1 to get better results, which would lead to around 60 000 evaluations with stochastics on stochastics if our method performs as well with that many points as with 5000 simulations (which is not guaranteed). We hope that this target will be reachable in the near future thanks to computing time improvements. It might however be complicated to use the same verification technique and to build convex hulls with that many points in general. However the polygon with 9 vertices in dimension 2 presented before is likely to work even better if many points are simulated, because the empirical quantiles of the radius will be closer and closer to the theoretical ones. Nevertheless, many questions still remain to be addressed: the models and the evaluation techniques should be improved (but this is not the topic of this paper); estimation problems might arise also because secondary simulations are used to evaluate the position of the company at time 1, and they should be controlled.

In the case where density level curves are unknown, some extreme value techniques could be used as a complement and provide some useful information. Note that here we would typically try to estimate level curves outside of which at least 4% or 6% of the points lie, whereas the methods developed for example by Einmahl (2010) for stable vectors are designed for cases where very few observations lie outside the level curve (and sometimes not even a single one). However, in some cases where heavy tails are involved, this theory could prove useful even if it would be at the limit of its applicability.

Another drawback of this two-step approach is that it implicitly requires the multi-dimensional risk process to be Markovian: one only relies on the current position of the risk drivers at time $t = 1$ to evaluate the company's net asset value at that time. It might be useful, at least for ORSA studies (Own Risk and Solvency Assessment) to have a better view on the speed and the direction in which markets and insurance risks are moving. Of course, some supplementary risk factors could be used to take into account some path characteristics, but this would be at the cost of dimension...

The main issue to be tackled is probably the curse of dimensionality: if one keeps all risk drivers of an insurance group, dimension could be around 200. In a partial internal model, dimension might be around 5 or 10 very often. In that case, dimension reduction is needed thanks to variable selection or other techniques. In fact, some risk factors might play a minor role, and consequently non-execution zones can become half-spaces or hyper-cylinders. We just highlight this phenomenon on a real-world portfolio with three risk drivers: asset risk, interest rate risk and property risk. Figures 12, 13 and 14 show the interest rate and the property risk drivers play a minor role. One way to get an idea of the importance of each risk driver is to use sensibilities, leading to 2-dimensional projected non-execution zones that are larger or smaller depending on the importance of each risk. In that case, only the asset risk is dominant, and consequently the 3-dimensional non execution zone is not far from a half-space (corresponding to low asset performances). With more risk drivers, it might be much more difficult to choose the main risk drivers and to control the error created by ignoring or using a proxy of the other risk drivers. We keep this for further research.

Our approach could also be combined with supplementary techniques. Liu and Staum (2010) use kriging techniques to choose the regions where the emphasis must be laid on after some first draws. This can be helpful to better take into account what




FIGURE 12. Case with 3 risk drivers: impact of asset risk and interest rate risk: red points are projections of the risk vectors with lower net asset value at time 1




FIGURE 13. Case with 3 risk drivers: impact of asset risk and property risk: red points are projections of the risk vectors with lower net asset value at time 1

we learn on the net asset value function as time goes by. However it might be too restrictive when risk factors are correctly assessed but dispersed in the outer sub-space (for example when both high and low interest rates are dangerous for the companies). Broadie et al. (2010) explain why it might be better to first roughly estimate the net




FIGURE 14. Case with 3 risk drivers: impact of interest rate risk and property risk: red points are projections of the risk vectors with lower net asset value at time 1

asset value at time 1 for each scenario with small secondary simulation numbers, and then to refine the estimations that are only needed. Gordy and Juneja (2008) suggest to consider first smaller number of insurance policies to get a first view of the net asset value function. This could be useful in cases where there is a large number of similar, homogeneous insurance policies. This is of course something that could be incorporated in our method to improve it. However, comparing these methods and their combinations is another story and we keep this for future research.

ACKNOWLEDGMENT

This work has been partially supported by the research chair *Management de la modélisation* sponsored by BNP Paribas Assurance, and by French Research National Agency (ANR) under the reference ANR-08-BLAN-0314-01.

REFERENCES

- Barnett, V. (1976) The ordering of multivariate data. *Journal of the Royal Statistical Society, Ser. A*, 139, 318-354.
- Broadie, M., Du, Y. and Moallemi, C. (2010) Efficient risk estimation via nested sequential simulation. To appear in *Management Science*.
- Chaouch, M., Gannoun, A., Saracco, J. (2008) Estimation de quantiles géométriques conditionnels et non conditionnels, *Société Française de Statistiques*
- Charpentier, A. and Oulidi, A. (2009) Estimating allocations for Value-at-Risk portfolio optimization, *Mathematical Methods of Operations Research*, Vol. 69, No. 3, 395-410.
- Chaudhuri, P. (1996). On a geometric notation of quantiles for multivariate data. *Journal of the American Statistical Association*, 91, 862-872.

Chaudhuri, P. and Sengupta, D. (1993) Sign tests in multidimension: inference based on the geometry of the data cloud. *Journal of the American Statistical Association*, 88, 1363-1370.

Devineau, L., Loisel, S. (2009a) Construction d'un algorithme d'accélération de la méthode des simulations dans les simulations pour le calcul du capital économique Solvabilité II, *Bulletin Français d'Actuariat*, No. 17, Vol. 10, 188-221.

Devineau, L., Loisel, S. (2009b) Risk aggregation in Solvency II: How to converge the approaches of the internal models and those of the standard formula?, *Bulletin Français d'Actuariat*, No. 18, Vol. 8, 107-145.

Einmahl, J. (2010) Estimating extreme quantile regions for multivariate regularly varying distributions, talk at the Journées de Statistique de la SFDS, Marseille.

Gordy, M. and Juneja, S. (2008) Nested Simulation in Portfolio Risk Measurement, Finance and Economics Discussion Series, Divisions of Research and Statistics and Monetary Affairs, Federal Reserve Board, Washington D.C.

Liu, M. and Staum, J. (2010) Stochastic Kriging for Efficient Nested Simulation of Expected Shortfall, forthcoming, *Journal of Risk*.

McNeil, A. (2010), Multivariate Stress Testing for Solvency II, talk at Risk Forum of the Institut Louis Bachelier, slides available at http://www.institutlouisbachelier.org/risk10/talk1_mcneil.pdf.

McNeil, A., Frey, R., Embrechts, P. (2005), *Quantitative Risk Management*, Princeton Series In Finance.

Nolan, J. (2005) Multivariate stable densities and distribution functions: general and elliptical case, Deutsche Bundesbank's 2005 Annual Fall Conference.

Rousseeuw, P., Ruts, I., and Tukey, J. (1999) The Bagplot: A Bivariate Boxplot, *The American Statistician*, Vol. 53.

Van der Vaart (1998), *Asymptotic Statistics*, Cambridge University Press.