Efficient computation of Solvency Capital Requirement using multilevel Monte Carlo methods

Alexandre Boumezoued
Adel Cherchali
Mathieu Truc

Many insurance companies are struggling to overcome the computational challenges involved in computing the Solvency Capital Requirement (SCR) under the Solvency II (SII) regime. Standard market approaches such as Least Square Monte Carlo (LSMC) and Replicating Portfolios (RP) can be difficult to calibrate and validate in practice but also lead to deficient outcomes if not calibrated properly. In this paper, we show how the Multilevel Monte Carlo (MLMC) method is a relevant alternative to compute risk-based capital requirements as it does not rely on any proxy assumptions.

The Solvency II directive is the prudential framework for insurers and reinsurers in Europe. The directive introduces the so-called Solvency Capital Requirement (SCR) to ensure that insurance companies are able to meet their financial claims. This framework introduces major innovation in the actuarial landscape such as market-consistent valuation of the balance sheet and risk-based capital requirements. To evaluate the SCR, the supervisory authority sets out two possible methodologies:

1. A "Standard Formula" approach, based on stress tests on several risk modules (e.g., interest rate, equity, mortality), and then an elliptical aggregation to derive an overall SCR.
2. An "Internal Model" approach that is based on a quantile of the one-year loss distribution of the insurer's portfolio at a 99.5% confidence level.

Our aim is to investigate numerical methods to compute the SCR using the internal model approach, which introduces major computational challenges. In a more general setting, the problem amounts to computing the probability of a large loss of a financial portfolio over a defined *risk horizon* τ. This task is particularly challenging in practice as complex insurance portfolios do not admit closed form solutions (for example because of embedded options or specific accounting rules). Hence, the valuation requires heavy Monte Carlo simulations. More formally, these types of problems involving simulations within simulations can be framed in the so-called *Nested Simulation* setting where outer scenarios are used to project the portfolio risk factors up to the risk horizon under the real-world probability, then inner simulations are necessary to compute the portfolio value conditionally on each primary scenario. This brute force approach is too time-consuming to be used in a real insurance business case.

In this paper, we introduce the MLMC methods developed by Giles (2008)[1] and their associated refinement strategies from Giles et al. (2019) .[2] This method relies on a smart allocation of a given computational budget between inner and outer scenarios spread across different *levels* to obtain an optimal trade-off between variance reduction and bias correction. This methodology exhibits major advantages for insurance companies as it does not rely on any proxy, which makes the approach easier to justify and validate. Indeed, the validation process for MLMC is relatively close to that of a full nested simulation approach, which would require justifying the chosen approach and retained parameters of the simulation (e.g., the number of outer and inner scenarios retained in a full nested exercise). This justification can require performing sensitivities to main parameters and assessing whether the estimated percentile value is stable, or more precisely measuring the optimality of the bias-variance trade-off. The reduction in the validation burden therefore comes from the "simulation" nature of the approach and the absence of any a priori assumption on the function that gives the response of Own Funds (OF) to risk factor outcomes. As such, there is no need to enter a process of proxy calibration, then validation and remediation, with any further calibration update if necessary, until a satisfactory replication is obtained. In addition, the MLMC algorithm can be efficiently parallelised and implemented in graphic cards (GPU) to further reduce the overall computational time. This makes the approach appealing under current and upcoming computational architectures leveraging capabilities from cloud services providers.

[1] Michael B Giles (2008). Multilevel Monte-Carlo path simulation. Operations research, 56(3):607–617

[2] Michael B. Giles & Abdul-Lateef Haji-Ali (2019). Multilevel nested simulation for efficient risk estimation. SIAM/ASA J. Uncertain. Quantif., 7(2):497–525.

This paper is organised as follows. Section 1 introduces the mathematical framework of the nested simulation approach and provides a precise definition of the quantile estimation problem, then Section 2 introduces the standard MLMC methods. Section 3 describes refinements that are relevant for SCR quantile estimation, while Section 4 illustrates the performance of the different algorithms on numerical experiments on a toy example. Finally, Section 5 shows applications to Internal Model SCR computation.

1. Nested simulations

The general problem is to estimate a risk measure of a financial portfolio over some future date τ called the risk horizon. We consider the general setting of Bauer et al. (2015).[3] Let V_0 be the current value of the insurance portfolio. The value of the portfolio at time τ can be expressed as a conditional expectation of future discounted cash flows under an Absence-of-Arbitrage (AOA) opportunity.

Let $(\Omega, \mathcal{F}, \mathcal{F}_{t\geq 0}, \mathbb{P})$ be a complete filter probability space modelling all possible market states and $\mathbb{P}$ the historical (real-world) probability measure. Let $(X_t)_{t\geq 0} = \left(X_t^1, \dots, X_t^d\right)_{t\geq 0}$ a d-dimensional Markov process that models the underlying risk factors of the portfolio. The filtration $\mathcal{F}_t$ represents all market information available up to time t and is generated by the underlying risk factors $\mathcal{F}_t = \sigma(\,(X_u)_{u\leq t})$. We also assume that a risk-neutral probability measure $\mathbb{Q}$ exists, under which discounted price processes are martingales. Let $Z \in \mathbb{R}$ be a one-dimensional random variable that represents the sum of future discounted cash flows generated by the portfolio. The (market) value of the portfolio at time τ is therefore given by:

$$V_\tau = \mathbb{E}^{\mathbb{Q}}[Z|\mathcal{F}_\tau]$$

The portfolio loss at time τ can be expressed as the change in the portfolio market value between 0 and τ:

$$L_\tau = V_0 - V_\tau = \mathbb{E}^{\mathbb{Q}}[V_0 - Z|\mathcal{F}_\tau]$$

To ensure that the company will remain solvent at time τ, we can compute the smallest amount x the company must hold today to have a small probability to make a large loss at the risk horizon (say less than α):

$$SCR = \inf\{x \in \mathbb{R}, \mathbb{P}(L_\tau \geq x) \leq \alpha\},$$

where the specifications underlying the SCR is $\tau = 1$ year and $\alpha = 0.5\%$; we note however that the proposed method is generic and can be applied to other specifications.

Denoting by F_{L_τ} the cumulative distribution function (cdf) of L_τ, the SCR can be found as the root of the following equation:

$$F_{L_\tau}(SCR) = 1 - \alpha$$

Hence the main problem is to estimate (both accurately and efficiently) the probability of a large loss:

$$I = \mathbb{P}(L_\tau \geq x)$$

This estimation problem can finally be written as a nested expectation problem:

$$I = \mathbb{E}^{\mathbb{P}}[g(\mathbb{E}^{\mathbb{Q}}[Y|\mathcal{F}_\tau])]$$

where

$$g(u) = \mathbb{I}_{u\geq x} \text{ and } Y = V_0 - Z$$

1.1 NESTED MONTE CARLO ESTIMATOR

The brute force estimation method is based on approximating the inner and outer expectation using independent Monte Carlo samples. The conditional inner expectation $\mathbb{E}^{\mathbb{Q}}[Y|\mathcal{F}_\tau]$ is estimated for a given realisation of the underlying risk factors $(x_0, \dots, x_\tau) \in (\mathbb{R}^d)^\tau$ by a standard Monte Carlo estimator with K inner simulations:

$$\hat{E}_K = \frac{1}{K}\sum_{k=1}^{K} Y_k$$

where $(Y_1, \dots, Y_K)$ is an i.i.d sample of the conditional distribution of Y given that $(X_0, \dots, X_\tau) = (x_0, \dots, x_\tau)$. The outer expectation is then approximated using the standard Monte Carlo estimator, using J outer simulations of the risk factors $\left(X_0^j, \dots, X_\tau^j\right)_{j=1,\dots,J}$:

$$\hat{I}_{J,K}^{Nested} = \frac{1}{J}\sum_{j=1}^{J} g\left(\hat{E}_{j,K}\right)$$

The procedure is illustrated in Figure 1.

FIGURE 1: NESTED MONTE CARLO METHOD

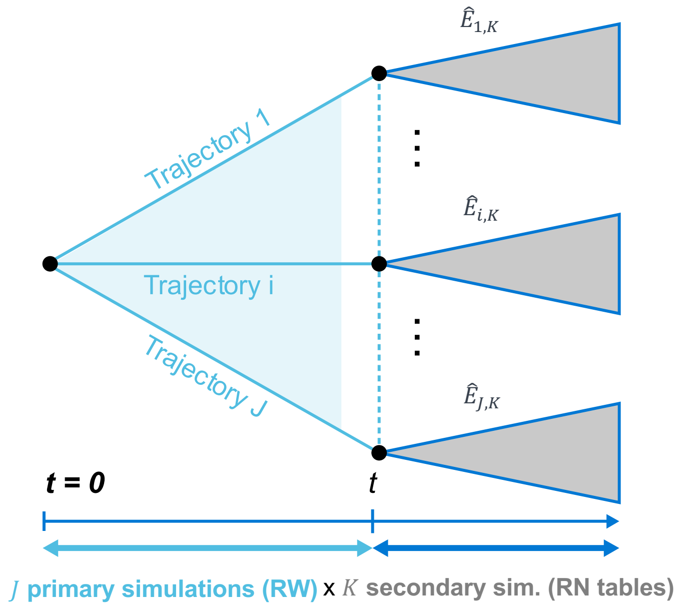

[3] Daniel Bauer & Hongjun Ha (2015). A least-squares Monte Carlo approach to the calculation of capital requirements. In World Risk and Insurance Economics Congress, Munich, Germany, August, pages 2–6.

1.2 COMPLEXITY ANALYSIS

Gordy et al. (2010)[4] and Hong et al. (2009)[5] analysed the best allocation strategies between inner and outer simulations to minimise the mean squared error (MSE) of the nested estimator for the indicator payoff $g = \mathbb{I}_{u\in[x,+\infty)}$ and to determine the best asymptotic complexity that can be achieved by the nested estimator.

For a given computational budget, the goal is to minimise the overall MSE of the estimator $\hat{I}_{J,K}^{Nested}$ using the bias-variance decomposition:

$$MSE\left(\hat{I}_{J,K}^{Nested}\right) = bias^2\left(\hat{I}_{J,K}^{Nested}\right) + Var\left(\hat{I}_{J,K}^{Nested}\right)$$

where:

$$bias\left(\hat{I}_{J,K}^{Nested}\right) = \mathbb{E}\left[\hat{I}_{J,K}^{Nested} - I\right]$$

The number of inner simulations K controls the level of bias, as typically increasing the number of inner simulations for a given Own Funds (OF) measurement based on risk-neutral simulations improved convergence to the true OF value. Also, the number of outer simulations J controls the level of variance (statistical error), as indeed the higher the number of real-world simulations (and OF value under each real-world simulation), the more precise the percentile estimate.

Based on asymptotic characterisation of the bias and variance, Broadie et al. (2011)[6] show the existence of an asymptotic optimal allocation $(J^\star, K^\star)$ that minimises the MSE. In particular, in order to achieve a root mean squared error (RMSE) in $O(\varepsilon)$, the optimal nested estimator requires $J^\star = O(\varepsilon^{-2})$ outer scenarios and $K^\star = O(\varepsilon^{-1})$ inner simulations, leading to an overall complexity in $O(\varepsilon^{-3})$.

As an order of magnitude, achieving a precision $\epsilon \sim 10^{-3}$ would require an optimal allocation of $\epsilon^{-2} \sim 10^6$ outer real-world simulations and $\epsilon^{-1} \sim 10^3$ inner risk-neutral scenarios, leading to a total simulation budget of $\epsilon^{-3} \sim 10^9$. The result above also shows that in general if we want to double the accuracy of the best possible nested estimator, we need $2^3 = 8$ times more simulations, which is either inefficient or out of reach in practice.

[4] Michael B. Gordy & Sandeep Juneja (2010). Nested simulation in portfolio risk measurement. Management Science, 56(10):1833–1848.

[5] L. Jeff Hong & Sandeep Juneja (2009). Estimating the mean of a non-linear function of conditional expectation. In Proceedings of the 2009 Winter Simulation Conference (WSC), pages 1223–1236. IEEE.

[6] Mark Broadie, Yiping Du, & Ciamac C. Moallemi (2011). Efficient risk estimation via nested sequential simulation. Management Science, 57(6):1172– 1194.

2. The MLMC method

2.1 OVERVIEW OF MLMC

Multilevel Monte Carlo (MLMC) methods have been successfully applied to compute nested expectation and reduce the overall complexity of the crude nested Monte Carlo estimator.

A MLMC method works as follows. Let $P_0, \ldots, P_L$ be a sequence of random variables approximating $P = g\left(\mathbb{E}^{\mathbb{Q}}[Y|\mathcal{F}_\tau]\right)$ with increasing accuracy and consequently increasing computational cost. The most accurate estimator of P is at the finer "level" denoted by L, leading to the approximation:

$$\mathbb{E}[P] \approx \mathbb{E}[P_L]$$

The error that comes from replacing P by P_L is the bias. The parameter L controls the depth of the bias correction.

The key ingredient of the methodology is that instead of estimation $\mathbb{E}[P_L]$ directly, we can expand it into a telescopic sum involving estimators at different levels l between 0 and L:

$$\mathbb{E}[P_L] = \mathbb{E}[P_0] + \sum_{l=1}^{L} \mathbb{E}[P_l - P_{l-1}]$$

It is then possible to approximate each expectation using a standard Monte Carlo procedure to get the MLMC estimator:

$$\hat{I}^{MLMC} = \frac{1}{J_0}\sum_{i_0=1}^{J_0} P_0^{(i_0)} + \sum_{l=1}^{L}\frac{1}{J_l}\sum_{i_l}^{J_l} P_l^{(i_l)} - P_{l-1}^{(i_l)}$$

The important point regarding the equation above is the flexibility in the method to allow for the number of outer simulations J_l to be different for different levels. The term at level 0 does most of the job in estimating I and it is the main term that contributes to reduce the variance of the estimator. The terms $P_l - P_{l-1}$ then aim at correcting the bias introduced by replacing P by P_L. To get an intuitive feeling of how the allocation of resources $l \mapsto J_l$ must be spread across levels, note that one should use high J_l for low levels l because P_l is cheap to compute at low levels. Then, if the sample $P_l - P_{l-1}$ can be made sufficiently negatively correlated, we are implicitly performing a variance reduction method, so we can safely decrease J_l without impacting the accuracy of the estimates in the region where $P_l - P_{l-1}$ are costly.

2.2 MLMC ESTIMATOR IN THE NESTED SIMULATION FRAMEWORK

In this section, we derive the MLMC estimator for nested expectation. Let $J = (J_l)_{l=0,\ldots,L}$ (or $K = (K_l)_{l=0,\ldots L}$) be the sequence modelling the number of outer simulations (or inner simulations) at each level. Following the intuitive reasoning of the previous section, $l \mapsto J_l$ must be decreasing while $l \mapsto K_l$ is increasing. A common choice is to consider the geometric progression on each level:

$$J_l = J_0 2^{-l}, K_l = K_0 2^l, l = 0, \ldots, L.$$

This leads to the nested version of the MLMC estimator:

$$\hat{I}^{MLMC} = \frac{1}{J_0}\sum_{i_0=1}^{J_0} g(\hat{E}_{i_0,K_0}) + \sum_{l=1}^{L}\frac{1}{J_l}\sum_{i_l=1}^{J_l} g(\hat{E}_{i_l,K_l}) - g(\hat{E}_{i_l,K_{l-1}})$$

We can observe that each level l has computational cost—

$$J_l \times K_l = J_0 \times K_0,$$

—while at any level l , the computation of $g(\hat{E}_{i_l,K_l})$ uses K_l inner scenarios and, because $K_{l-1} = \frac{K_l}{2}$, half of the sample is "thrown away" to compute $\hat{E}_{i_l,K_{l-1}}$. We can therefore "recycle" the second half sample to perform variance reduction.
This observation leads to the so-called antithetic version of the MLMC estimator:

$$\hat{I}^{MLMC} = \frac{1}{J_0}\sum_{i_0=1}^{J_0} g(\hat{E}_{i_0,K_0}) + \sum_{l=1}^{L}\frac{1}{J_l}\sum_{i_l=1}^{J_l}\left\{g(\hat{E}_{i_l,K_l}) - \frac{g(\hat{E}_{i_l,K_{l-1}}) + g(\hat{E}'_{i_l,K_{l-1}})}{2}\right\}$$

where $\hat{E}'_{i_l,K_{l-1}}$ is the empirical mean over the second half of the sample at next level:

$$\hat{E}'_{i_l,K_{l-1}} = \frac{1}{K_{l-1}}\sum_{k=\left\lceil\frac{K_{l-1}}{2}\right\rceil}^{K_l} Y_{i_l,k}$$

Giles (2008) shows that a MLMC estimator of this type can reduce the computational cost from $O(\varepsilon^{-3})$ to $O(\varepsilon^{-2})$ depending on the regularity of the payoff function.

Under the latter complexity case, if we want to double the accuracy of the estimation, one now only needs $2^2 = 4$ times more simulations, which is a significant reduction with regard to the crude nested estimator performance. In addition, this result roughly states that antithetic type MLMC estimator can be reduced to an unbiased Monte Carlo estimation, as if the conditional expectation $\mathbb{E}[Y|X]$ were known in closed form. Indeed, it is known that the standard Monte Carlo method converges in $O(1/\sqrt{N})$, hence, to reach an accuracy ε within the standard Monte Carlo framework, one needs roughly $N = O(\varepsilon^{-2})$ scenarios.

3. Adaptive MLMC

Unfortunately for the problem at hand (quantile estimation), the payoff function $g(u) = \mathbb{I}_{u\geq x}$ is not smooth enough to achieve the $O(\varepsilon^{-2})$, and the standard MLMC estimator complexity is in fact $O\left(\varepsilon^{-\frac{5}{2}}\right)$ in this case. Hence, further improvement is necessary to reach this optimal complexity.

Giles et al. (2019)[7] reduce the complexity to $O(\varepsilon^{-2}\log(\varepsilon)^2)$ using an adaptive strategy that uses a random allocation of inner scenarios across levels instead of a geometric progression based on the idea of Broadie et al. (2011).
The idea works as follows:

Let us consider a payoff function $g(u) = \mathbb{I}_{u\geq x}$ and that we want to estimate the following quantity:

$$I = \mathbb{P}(\mathbb{E}^{\mathbb{Q}}[Y|X] \geq x) = \mathbb{E}[g(\mathbb{E}^{\mathbb{Q}}[Y|X])]$$

Consider two outer scenarios X^1, X^2 such that:

Case 1: The true value of the conditional expectation is close to the loss threshold ($\mathbb{E}^{\mathbb{Q}}[Y|X^1] \approx x$), hence the area where the indicator function is not continuous (values 0 or 1).

Case 2: The true conditional expectation is very far from the singularity ($\mathbb{E}^{\mathbb{Q}}[Y|X^2] \gg x$ or $\mathbb{E}^{\mathbb{Q}}[Y|X^2] \ll x$).

The graph in Figure 2 illustrates the typical situation.

FIGURE 2: STOCHASTIC ALLOCATION OF INNER SCENARIOS

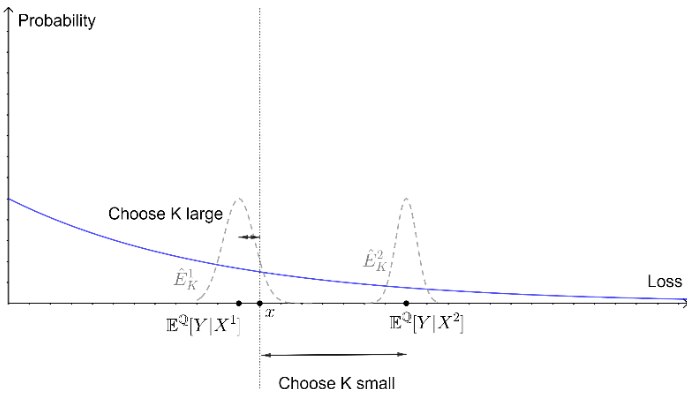

In Case 2, a uniform strategy (J, K) is clearly not optimal, as in that case we can afford only a very rough estimate of $\mathbb{E}^{\mathbb{Q}}[Y|X^2]$ using only a few samples. The estimation of $g(\mathbb{E}^{\mathbb{Q}}[Y|X^2])$ would be 1 because we are very far from the threshold of the payoff. However, in Case 1 ($\mathbb{E}^{\mathbb{Q}}[Y|X^1] \approx x$), there is a high risk of *misclassification* and many more inner scenarios must be allocated to get a good estimate of $g(\mathbb{E}^{\mathbb{Q}}[Y|X^1])$. Thus, the idea is to identify these high-risk primary scenarios X^i by considering those maximising the probability of a change of sign if we just add one more scenario. The related optimisation problem is:

$$i^\star = argmax_{i=1,\dots,J}\ \mathbb{P}(\hat{E}_{i,K_i+1} < x | \hat{E}_{i,K_i} > x)$$

From Bienaymé-Tchebetchev inequality we get:

$$\mathbb{P}(\hat{E}_{i,K_i+1} < x | \hat{E}_{i,K_i} > x) \leq \frac{Var(Y|X^i)}{K_i^2 |\hat{E}_{i,K_i} - x|^2}$$

[7] Giles & Haji-Ali (2019), op cit.

This result suggests that we should allocate more inner scenarios if we are close to the threshold of the payoff function x (i.e., $|\hat{E}_{i,K_i} - x|$ small in the denominator), with a current low number of inner simulations K_i and a high variance $Var(Y|X^i)$.

Therefore, to be sure that the estimate $g(\hat{E}_{i,K_i}) \approx g(\mathbb{E}^{\mathbb{Q}}[Y|X^i])$ is accurate, i.e., that the probability to make a change of sign is small (less than ε), we should take K such that

$$K \geq \frac{\sigma}{d}\varepsilon^{-\frac{1}{2}}$$

with $\sigma = \sqrt{Var(Y|X)}$, $d = |\mathbb{E}^{\mathbb{Q}}[Y|X] - x|$ and impose a cap on the number of maximum samples to use to keep the complexity under control:

$$K = \min\left\{O(\varepsilon^{-1}), \frac{\sigma}{d}\varepsilon^{-\frac{1}{2}}\right\}$$

Broadie et al. (2011) showed that this stochastic allocation of resources improved the complexity of the crude nested Monte Carlo estimator initially of $O(\varepsilon^{-3})$ down to $O\left(\varepsilon^{-\frac{5}{2}}\right)$ when such adaptive strategy is considered.

How can this strategy be leveraged to further improve the MLMC estimator?

Firstly, observe that an alternative way to construct a stochastic rule for K is motivated by the central limit theorem because the confidence interval for the estimator $\hat{E}_{x,K}$ for a given $X = x$ is of the form:

$$\hat{E}_{x,K} \in \left[E[Y|X = x] \pm C\sqrt{\frac{Var(Y|X = x)}{K}}\right]$$

This leads to the following alternative rule:

$$K \sim \frac{\sigma^2}{d^2}C^2$$

In the context of MLMC, the idea is to introduce a parameter $r \in (1,2)$ and set the number of inner scenarios as the middle ground between these two allocation rules; the final formula to set the number of inner scenarios at level l is

$$K_l = K_0 4^l \max\left\{2^{-l}, \min\left\{1, \left(C^{-1}K_0^{\frac{1}{2}}2^l\frac{d}{\sigma}\right)^{-r}\right\}\right\}$$

with the following interpretation:

- The number of samples at each level l must be comprised in $[K_0 2^l, K_0 4^l]$ so that at most we generate 2^l times more inner scenarios.
- The introduction of a parameter $r \in (1,2)$ that is the middle ground power between the two-allocation rule introduced before that controls the number of samples. As r increases, the number of inner samples decreases.
- The constant C corresponds to the confidence bound constant.

4. Numerical experiments

In this section, we illustrate the theoretical results on a simple example where the conditional expectation and the value at risk are known in closed form.

4.1 SETTING

We consider a market model with one asset following a Black-Scholes diffusion:

$$\frac{dS_t}{S_t} = \mu dt + \sigma dW_t$$

We consider a portfolio comprising a single put option of maturity T with strike K. The price of the put option at time $t \leq T$ is known in closed form:

$$\begin{aligned} P_t &= \mathbb{E}^{\mathbb{Q}}\left[e^{-r(T-t)}(K - S_T)_+ | S_t\right] \\ &= Ke^{-r(T-t)}\mathcal{N}(-d_-) - S_t\mathcal{N}(-d_+) \end{aligned}$$

with:

$$d_{\pm} = \frac{\log\left(\frac{S_t}{Ke^{-r(T-t)}}\right) \pm \frac{\sigma^2}{2}(T-t)}{\sigma\sqrt{T-t}}$$

And $\mathcal{N}(.)$ the cdf of the standard normal distribution.

This example, although relatively simple, is relevant in the insurance setting as it represents the cost of embedded options: An insurer proposing a minimum guaranteed rate r_G amounts to selling a put option with strike $K = r_G$. Hence the probability to make a large loss because of the cost of such embedded options is the building block of computations of interest in practice.

The portfolio loss at time τ is given by:

$$\begin{aligned} L_\tau &= P_0 - P_\tau \\ &= \mathbb{E}^{\mathbb{Q}}\left[P_0 - e^{-r(T-\tau)}(K - S_T)_+ | S_\tau\right] \end{aligned}$$

The value at risk (VaR) at level α denoted $VaR_\alpha(L_\tau)$ is:

$$Var_\alpha(L_\tau) = P_0 - Var_\alpha(P_\tau)$$

where:

$$Var_\alpha(P_\tau) = Ke^{-r(T-\tau)}\mathcal{N}(-d_-(x_\alpha)) - x_\alpha\mathcal{N}\left(-d_+(x_\alpha)\right)$$

and

$$x_\alpha = S_0 e^{\left(r - \frac{\sigma^2}{2}\right)\tau + \sigma\sqrt{\tau}\mathcal{N}^{-1}(1-\alpha)}$$

We want to compare the performance of the MLMC algorithms (original geometric version, as well as adaptive) with the crude nested Monte Carlo estimator in these settings. The goal is to compute the following quantity:

$$Var_\alpha(L_\tau) = F_{L_\tau}^{-1}(1-\alpha)$$

where $F_{L_\tau}^{-1}(.)$ is the inverse cdf of the one-year loss distribution at time τ. Therefore, the VaR problem can be viewed as a root search problem:

$$Find\ x\ s.t\ \mathbb{P}(L_\tau \geq x) = \alpha$$

And we are back to the problem of estimating the probability of a large loss in a nested framework:

$$I = \mathbb{P}(L_\tau \geq x) = \mathbb{E}^{\mathbb{P}}[g(\mathbb{E}^{\mathbb{Q}}[Y|X])]$$

with $X = S_\tau$, $Y = P_0 - e^{-r(T-\tau)}(K - S_T)_+$.

4.2 NUMERICAL RESULTS

In our numerical experiments, we considered the initial stock price $S_0 = 100$, the volatility $\sigma = 30\%$, the option maturity $T = 5$ years, and the risk horizon $\tau = 1$ year.

In our tests, we deal with two cases:

- At-the-money (ATM): $S_0 = K = 100$.
- Deep-in-the-money (ITM): $S_0 \ll K = 200$. In this case, the put price is close to the forward price struck at K because the payoff function is close to linear, hence we are in a case similar to a rare event estimate because the indicator payoff $1 - \mathbb{I}_{\mathbb{E}[Y|X] \geq x} \approx 0$ most of the time.

Asymptotic complexity for the probability of a large loss

We have drawn in the following figure the root mean squared error (RMSE) as a function of the computational cost. We generated N_{batch} = 150 replications of the algorithms to derive the empirical cost.

The computational cost is defined by the following quantities:

- Nested Monte Carlo: $Cost = J \times K$.
- Multilevel geometric: $Cost = \sum_{l=1}^{K_l} J_l K_l$.
- Multilevel adaptive: $Cost = \sum_{l=1}^{K_l} J_l \widehat{K}_l$, with $\widehat{K}_l = \frac{1}{N_{batch}} \sum_{k=1}^{N_{batch}} K_l^k$ the average inner scenarios per level.

In Figures 4 and 6, we observe roughly a behaviour in ε^{-3} for the nested estimator (blue dotted line). In Figure 3 and 4, we observe a complexity in order of $\varepsilon^{-\frac{5}{2}}$ for the MLMC geometric estimator (blue dotted line), which is consistent with the theoretical result of Giles et al. We also observe that the adaptive version (orange dotted line in Figures 3 and 4) outperforms the standard MLMC approaches because the slope of the line is smaller than the geometric MLMC. This is consistent with the theoretical asymptotic of the adaptive approach which is in $\varepsilon^{-2} log(\varepsilon)^2$. The adaptive version of the MLMC algorithm is particularly well suited for rare event estimation (ITM case, Figure 5) because we are performing importance sampling by construction, which leads to an improvement in term of variance reduction compared to the geometric version.

FIGURE 3: EMPIRICAL RMSE OF THE MLMC ESTIMATORS AS A FUNCTION OF THE COST IN LOG-SCALE (ATM CASE)

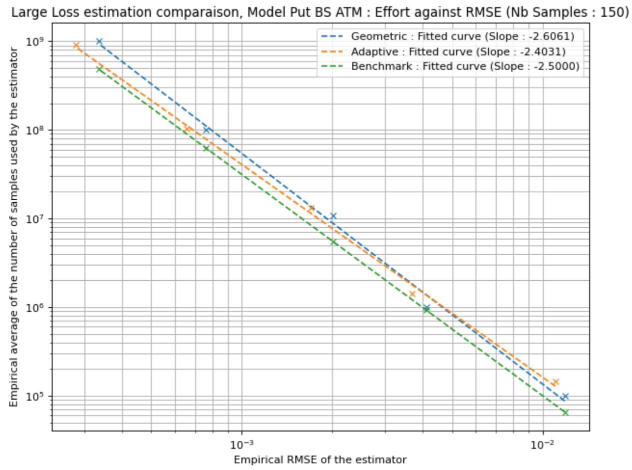

FIGURE 4: EMPIRICAL RMSE OF THE NESTED ESTIMATOR AS A FUNCTION OF THE COST IN LOG-SCALE (ATM CASE)

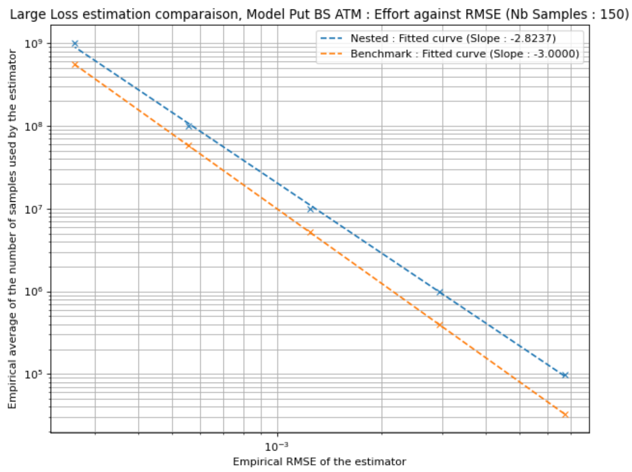

FIGURE 5: EMPIRICAL RMSE OF THE MLMC ESTIMATORS AS A FUNCTION OF THE COST IN LOG-SCALE (ITM CASE)

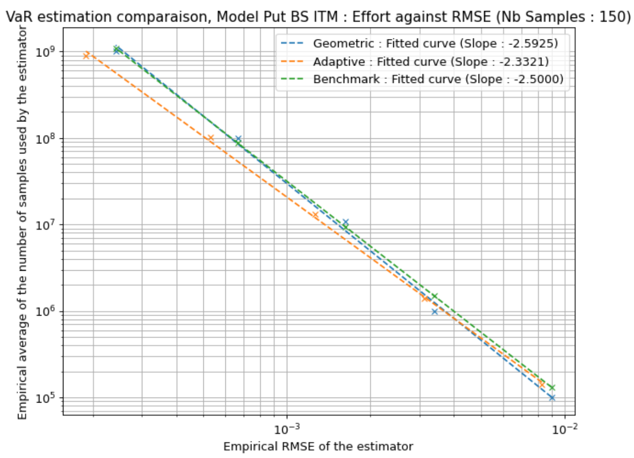

FIGURE 6: EMPIRICAL RMSE OF THE NESTED ESTIMATOR AS A FUNCTION OF THE COST IN LOG-SCALE (ITM CASE)

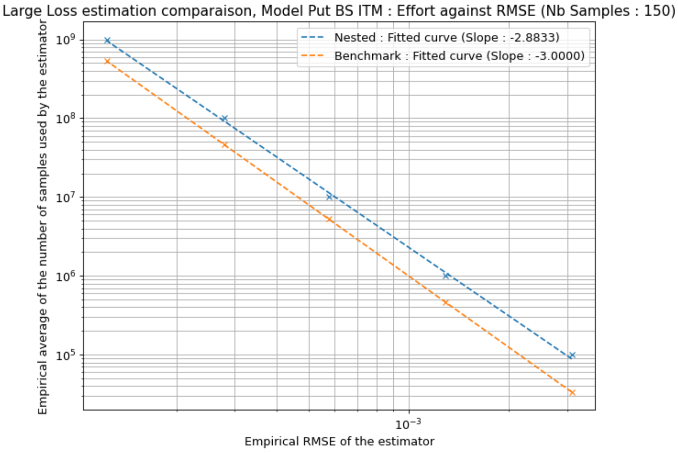

Asymptotic complexity for the VaR

We now present the asymptotic result for the quantile estimates. In this case, we use a numerical inversion of the empirical survival function $x \mapsto 1 - \hat{F}_N(x)$ obtained with each algorithm (nested, MLMC geometric, MLMC adaptive). This numerical inversion introduces an additional error compared to the computation of the probability of a large loss I. The plot in Figure 7 displays the empirical survival function of the MLMC algorithms (Figures 7 and 8) and compared the estimated VaR with the exact value at risk.

FIGURE 7: SURVIVAL FUNCTION AND VAR ESTIMATION USING THE GEOMETRIC MLMC ALGORITHM

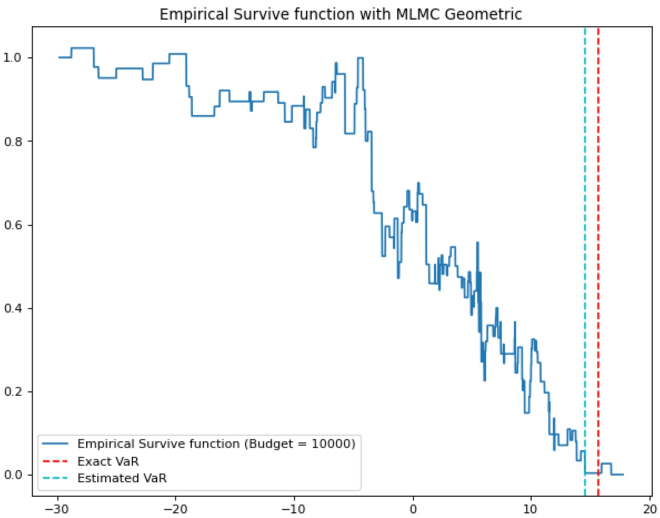

FIGURE 8: SURVIVAL FUNCTION AND VAR ESTIMATION USING THE ADAPTIVE MLMC ALGORITHM

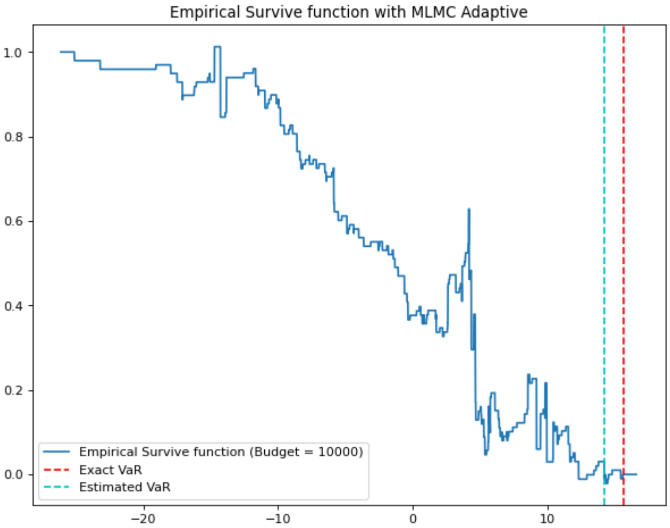

In Figures 9 and 10, we observe that the asymptotic complexity results for the VaR are globally in line with the theoretical result shown in the previous section. However, let us observe that the asymptotic complexity for the MLMC methods is slightly deteriorated because of the additional error due to the numerical inversion of the empirical cdf. In any case, the adaptive version of the MLMC algorithm outperforms the standard MLMC approach.

FIGURE 9: EMPIRICAL RMSE OF THE NESTED AND MLMC ESTIMATORS AS A FUNCTION OF THE COST IN LOG-SCALE (ATM CASE)

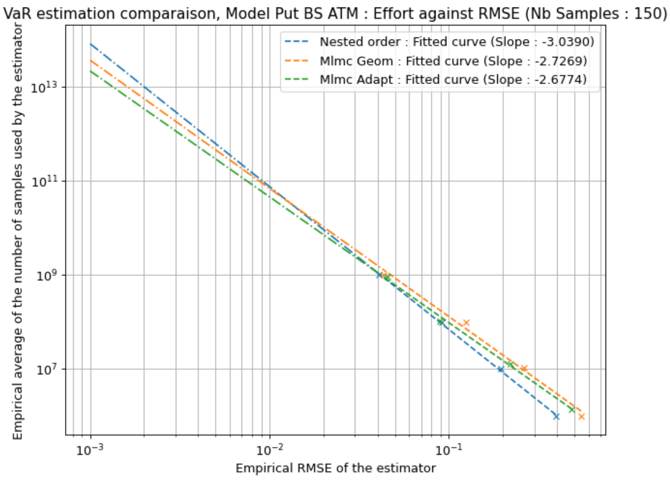

FIGURE 10: EMPIRICAL RMSE OF THE NESTED AND MLMC ESTIMATORS AS A FUNCTION OF THE COST IN LOG-SCALE (ITM CASE)

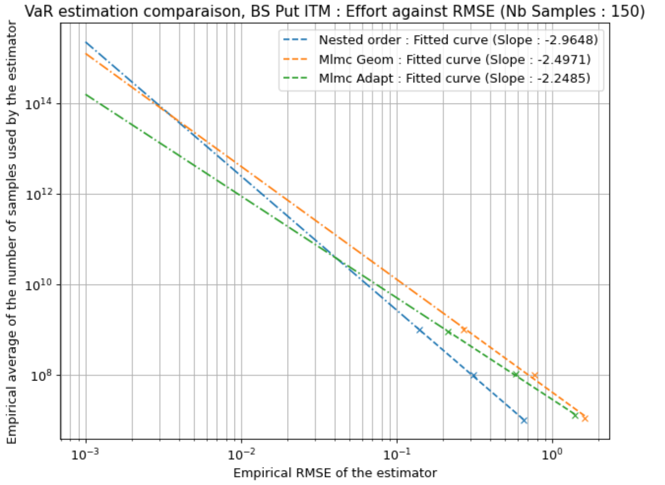

5. Application to Internal Models

In this section, we illustrate the performance of the MLMC methods on a realistic example for insurance applications. We consider the case of asset and liability management (ALM) for a life insurance portfolio and are interested in the derivation of the Internal Model SCR. Here, we use the framework developed by Floryszczak et al. (2016)[8]—a description of the main steps of the model and the main notation can be found in the Appendix section. We refer to this paper for further details.

We now present numerical results for the calculation of the SCR in the ALM model. We use the set of parameters shown in Figures 11 and 12 for the ALM model and the asset model; they are described in the Appendix.

FIGURE 11: ALM PARAMETERS

r_G	PSR	M	H	L_0	V_0^b	c_0	r_{SL}	β	r_{SL}	ψ	T	x
1%	80%	10Y	105	120	96	30	50	40%	5%	5%	30	75

FIGURE 12: FINANCIAL MARKET MODEL PARAMETERS

S_0	r_0	σ_S	σ_r	κ	θ	ρ
45	3%	30%	3%	20%	5%	50%

We focus on the calculation of the SCR (quantile formulation) at a one-year horizon. In Figure 13, we have drawn the RMSE of the nested estimator $\widehat{SCR}_{Nested}$, standard MLMC estimator $\widehat{SCR}_{MLMC,geom}$, and adaptive MLMC estimator $\widehat{SCR}_{MLMC,adaptive}$. To derive the RMSE of the different estimators, as no closed formula is available in this framework, we rely on a full nested Monte Carlo procedure based on a fixed computational budget $\Gamma = 10^8$ sample path to approximate the true value of the SCR. To compute the RMSE of the different estimators we produce $N_{batch} = 50$ simulations $\left(\widehat{SCR}_j\right)_{j=1,\dots,N_{batch}}$ and compute the empirical RMSE given by:

$$RMSE = \sqrt{\frac{1}{N_{batch}} \sum_{j=1}^{N_{batch}} \left(\widehat{SCR}_j - SCR\right)^2}$$

We plot the empirical RMSE's of the different estimators as a function of the computational cost in log-scale.

FIGURE 13: EMPIRICAL RMSE OF THE NESTED ESTIMATOR AS A FUNCTION OF THE COST IN LOG-SCALE

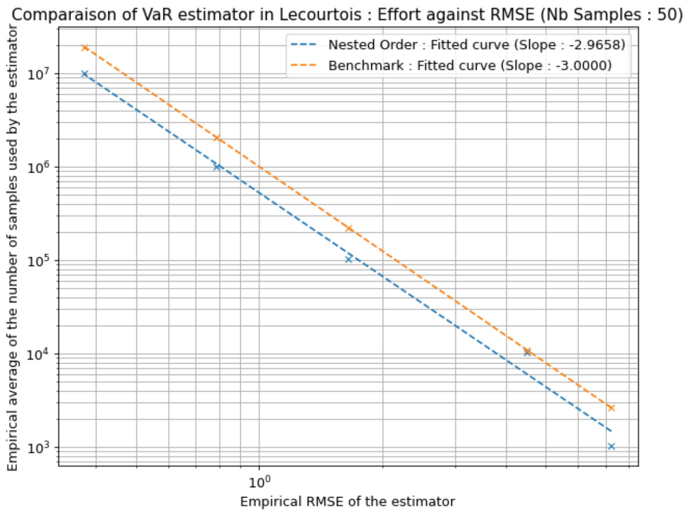

[8] Floryszczak, Anthony, Olivier Le Courtois, and Mohamed Majri (2016). Inside the Solvency 2 black box: Net asset values and solvency capital requirements with a least-squares Monte-Carlo approach. Insurance: Mathematics and Economics 71: 15-26.

FIGURE 14: EMPIRICAL RMSE OF THE MLMC ESTIMATORS AS A FUNCTION OF THE COST IN LOG-SCALE

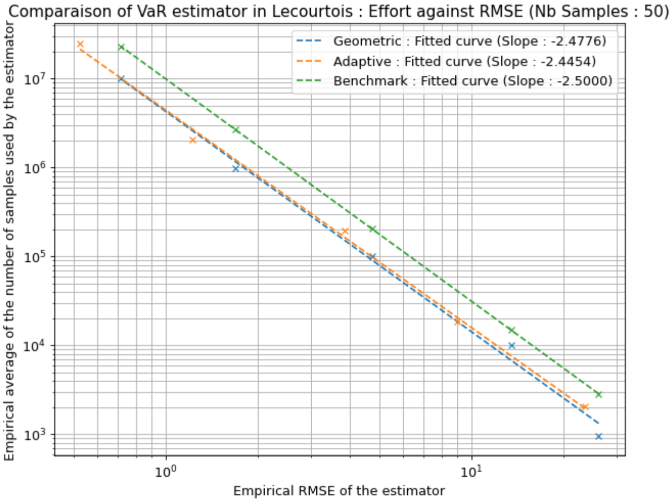

In Figure 13, we observe that the behaviour of the nested estimator is roughly in $O(\varepsilon^{-3})$ while MLMC methods are approximately in order $O\left(\varepsilon^{-\frac{5}{2}}\right)$, which is in line with the results of Giles et al. In this case, adaptive MLMC slightly outperforms the standard MLMC estimator but does not attain the optimal complexity $O(\varepsilon^{-2})$. Nevertheless, it represents a significant improvement compared to a crude nested estimator.

Conclusion

In this paper, we have shown that MLMC methods are a relevant alternative to the proxy modelling approaches that are commonly used in capital requirement estimation within the insurance sector. This approach represents a paradigm shift in that it smartly allocates computational resources to correct the inner bias of the nested procedure instead of relying on a proxy. From the operational point of view, this method eases the overall validation process of the insurance company.
In proxy modelling, explanatory variables must be carefully selected and the proxy function must be properly calibrated and validated. With MLMC methods, this complex process is bypassed because no proxy is involved in the estimation process. To our knowledge, this is the first time an adaptive MLMC algorithm has been implemented for insurance applications. This adaptive version of the algorithm outperforms the standard MLMC estimator and can be efficiently parallelised and implemented on graphic cards to obtain a significant performance boost. With the recent development of cloud computing and improved computational capabilities within the insurance industry, MLMC methods might be one high-potential solution to unlock the nested simulation problem and estimate the SCR efficiently.

Appendix

In this appendix section, we describe the main steps of the ALM framework developed in by Floryszczak et al. (2016).[9]

Savings contract characteristics

We consider an insurance company that handles savings contracts. The policyholder makes a deposit (their savings) and the insurance company guarantees a minimal earning (minimal guaranteed rate r_G) each year. This amount is then invested in the financial market and the policyholder is granted an additional bonus called Profit-Sharing, corresponding to a proportion $PSR \in [0,1]$ of the gain of the financial portfolio. The contract terminates upon death or surrender of the policyholder.

The initial deposit of the policyholder is invested in stocks, sovereign bonds and cash (i.e., deposited in a bank account). A part of the capital is invested in coupon-bearing bonds of maturity M with nominal H delivering each year coupon payment c_b with market-value:

$$V_t = \sum_{i=1}^{M} c_b H\, P(t, t+i) + H\, P(t, t+M)$$

The initial value of the portfolio is given by:

$$A_0 = V_0 + S_0 + C_0$$

where S_0 is the stock price at time 0 and C_0 the initial cash level.

Surrender outflow and cash dynamics

Outflows occur upon death or lapsation of the policyholder. Let q_{x+t-1} be the probability that a policyholder with age $x+t-1$ at time t die next year. Let $r_L(\Delta)$ be the lapse probability between two dates. The cash outflow is given by:

$$F_t = L_{t-1}\left(q_{x+t-1} + \psi r_L(\Delta)\right)$$

where $(L_t)_{t\geq 0}$ is the liability book-value (initial deposit and accrued credited rate) and ψ a penalty in case of surrender.

The lapse rate is modelled as a parabolic function:

$$r_L(\Delta) = r_{SL} + \mathbb{I}_{\Delta>0} \min\{\alpha\Delta^2, \beta\}$$

where $\Delta = r - r_C$ is the difference between the market rate r and the crediting rate r_C. The parameter r_{SL} quantifies structural surrenders, α is the speed rate at which a policyholder lapses their contract if the rate proposed by the insurance company is too low compared to the market rate and β is a maximal (cap) surrender rate.

[9] Ibid.

Cash dynamics

The cash account is a deposit on a bank account that yields interest rate r_t. The coupons from investment are deposited on this account and outflows are paid with cash in priority:

$$c_t = c_{t-1}e^{\int_{t-1}^{t} r_s ds} + Hc_b - F_t$$

If there is not enough cash, bonds and stocks are sold on the market to provide the necessary liquidity.

Asset management

When bonds mature, all cash except a security amount is invested in new bonds of the same maturity. At time M, the market value of the bond satisfies:

$$V_M^b = (c_M - c_0) = \sum_{i=1}^{M} c_b H^\star P(t, t+i) + H^\star P(t, t+M)$$

where c_0 is the security amount and $H^\star$ the nominal value that solves this equation.

The asset return between $[t-1, t)$ is measured by the log-variation of the price:

$$R_t^a = log\left(\frac{A_t}{A_{t-1}}\right)$$
$$= \log\left(\frac{V_t^b + S_t + c_{t-1}e^{\int_{t-1}^{t} r_s ds} + Hc_b}{V_{t-1}^b + S_{t-1} + c_{t-1}}\right)$$

Liability dynamics

We distinguish two types of liability to compute the debt toward the policyholder:

- **Market value of liability:** The maximum between the guaranteed amount and the profit-sharing, taking into account the cash outflows:

$$\mathcal{L}_t = \mathcal{L}_{t-} - F_t$$

where :

$$\mathcal{L}_{t-} = max\{\mathcal{L}_{t-1}(1 + PSR \times R_t^a), L_{t-1}e^{r_G}\}$$

- **Liability book value:** This quantity corresponds to policyholders' savings updated with the crediting-rate $r_C(t)$:

$$L_t = L_{t-1}\big(1 + r_C(t)\big)$$

where the crediting rate takes the following form:

$$r_C(t) = \underbrace{e^{r_G}}_{minimum\ guaranteed\ rate} + \underbrace{PSR\left(\frac{\mathcal{L}_{t-}}{L_{t-1}} - e^{r_G} - 1\right)_+}_{additionnal\ profit\ sharing} - 1$$

Solvency II balance sheet

The balance sheet of the insurance company is divided into two groups:

- **Asset:** The market-value A_t of the asset.
- **Liability:** The liability side is decomposed into two groups. Firstly the so-called *Best-Estimates of Liabilities* (BEL) corresponds to the discounted value of future cash outflows. This is the estimated debt of the company that sells the insurance contract:

$$BEL_t = \mathbb{E}^{\mathbb{Q}}\left[\sum_{u=t}^{T} e^{-\int_t^u r_s ds} F_u | \mathcal{F}_t\right]$$

where T is the maturity of the insurance contract. The Own Fund of the company (also called Net-Asset-Values or NAV) is the amount that remains when the company has paid its debt. A negative NAV is a situation of insolvency:

$$NAV_t = A_t - BEL_t$$

The Solvency Capital Requirement (SCR) is the minimum amount $x^\star$ that must be added to the current Own Fund to avoid insolvency in the next year with a high confidence level $1-\alpha$:

$$SCR_{t+1} = \inf\left\{x \in \mathbb{R}, \mathbb{P}\left(NAV_{t+1} + xe^{\int_t^{t+1} r_s ds} \geq 0 | \mathcal{F}_t\right) \geq 1 - \alpha\right\}$$

Financial market model

The fund managers invest policyholder deposits in stocks, cash and bonds; therefore we need stochastic models to generate the stock index and the level of interest rate. For the purpose of illustration, we consider Black-Scholes-Vasicek dynamics:

$$\frac{dS_t}{S_t} = r_t dt + \sigma_S dW_t^S$$
$$dr_t = \kappa(\theta - r_t)dt + \sigma_r\left(\rho dW_t^S + \sqrt{1-\rho^2}dW_t^r\right)$$

References

- Alfonsi, A., Cherchali, A., & Acevedo, J. A. I. (2021). Multilevel Monte-Carlo for computing the SCR with the standard formula and other stress tests. Insurance: Mathematics and Economics, 100, 234-260.
- Alfonsi, A., Cherchali, A., & Infante Acevedo, J. A. (2020). A synthetic model for asset-liability management in life insurance, and analysis of the SCR with the standard formula. European Actuarial Journal, 10(2), 457-498.
- Daniel Bauer and Hongjun Ha. A least-squares monte carlo approach to the calculation of capital requirements. In World Risk and Insurance Economics Congress, Munich, Germany, August, pages 2–6, 2015.
- Bourgey, F., De Marco, S., Gobet, E., & Zhou, A. (2020). Multilevel Monte Carlo methods and lower–upper bounds in initial margin computations. Monte Carlo Methods and Applications, 26(2), 131-161.
- Floryszczak, Anthony, Olivier Le Courtois, and Mohamed Majri. "Inside the Solvency 2 black box: net asset values and solvency capital requirements with a least-squares Monte-Carlo approach." Insurance: Mathematics and Economics 71 (2016): 15-26.
- Michael B Giles. Multilevel Monte-Carlo path simulation. Operations research, 56(3):607–617, 2008.
- Michael B. Giles and Abdul-Lateef Haji-Ali. Multilevel nested simulation for efficient risk estimation. SIAM/ASA J. Uncertain. Quantif., 7(2):497–525, 2019.
- Giorgi, D., & Lemaire, V. (2020). Weak error for nested multilevel Monte Carlo. Methodology and Computing in Applied Probability, 22(3), 1325-1348.Risa S. (2001).
- Michael B. Gordy and Sandeep Juneja. Nested simulation in portfolio risk measurement. Management Science, 56(10):1833–1848, 2010.
- Frank E. Harrell and C. E. Davis. “A new distribution-free quantile estimator”. In: Biometrika 69.3 (Dec. 1982), pp. 635–640. issn: 0006-3444. doi: 10.1093/biomet/69.3.635.
- L Jeff Hong and Sandeep Juneja. Estimating the mean of a non-linear function of conditional expectation. In Proceedings of the 2009 Winter Simulation Conference (WSC), pages 1223–1236. IEEE, 2009.
- Lemaire, V., & Pagès, G. (2017). Multilevel Richardson–Romberg extrapolation. Bernoulli, 23(4A), 2643-2692.

CONTACT

Alexandre Boumezoued
alexandre.boumezoued@milliman.com

Adel Cherchali
adel.cherchali@milliman.com

Mathieu Truc
mathieu.truc@milliman.com