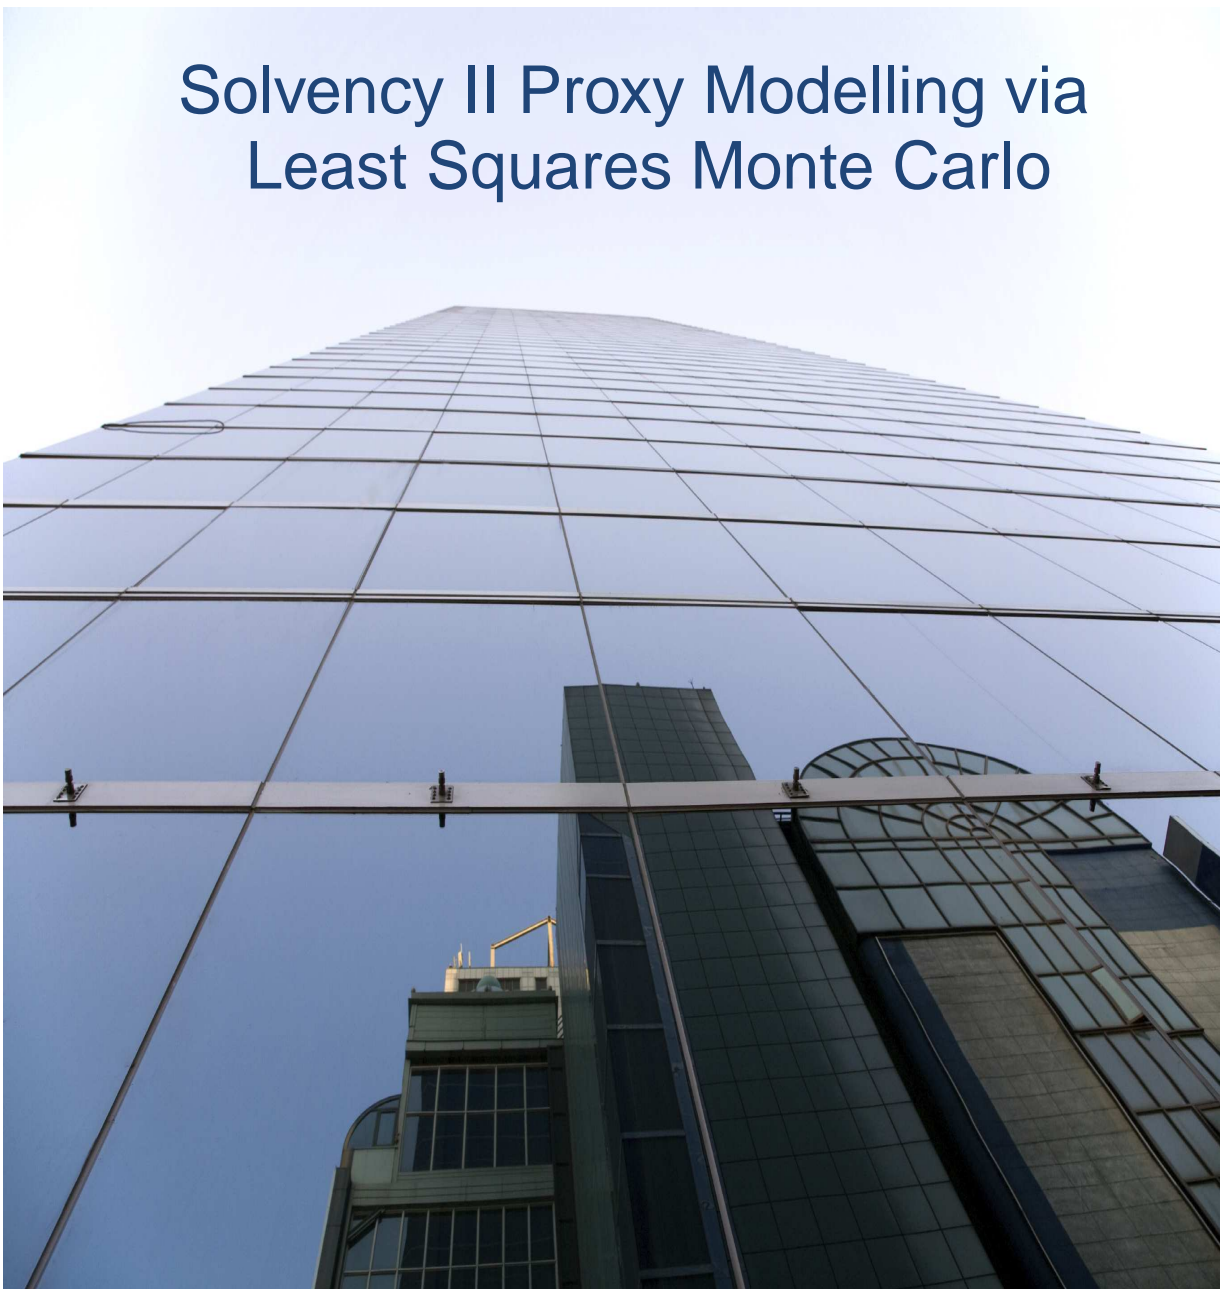


Mario Hörig, Michael Leitschkis

January 2012



Solvency II Proxy Modelling via Least Squares Monte Carlo



Executive Summary

For European life insurers which have decided to implement Solvency II on an internal model basis it is necessary to consider two stochastic dimensions: a so called “outer” real-world one to capture the probability distribution of different outcomes for the economic balance sheet over a one year time horizon and an “inner” risk neutral one in order to value assets and liabilities along each of these real world paths. Such a calculation would require a nested stochastic approach, which for many companies is very difficult or simply impractical due to the number of runs necessary.

This has led companies to look for proxy modelling techniques which can estimate the full approach in more acceptable run times. Methods used to date include replicating portfolios and curve fitting, but neither of these techniques is straightforward and when there are complex interactions between assets and liabilities, there may be difficulty getting an adequate level of fit. Recently a lot of attention has focused on another technique, Least Square Monte Carlo or LSMC. In the LSMC method we do not use thousands, but just a small number - such as 10 - of inner valuation scenarios for each outer one. This gives a very inaccurate valuation, but by carrying out a regression we can arrive at a very good estimate of the precise calculation we would get from a full nested stochastic approach.

Via LSMC, we can obtain very accurate results at a fast run time. The accuracy of LSMC results can be verified in a practically robust and statistically sound way. The LSMC method requires much less manual intervention than some of the alternatives and can give valuable economic insights about the interplay of different risk drivers. In this paper we explain the process used to carry out the LSMC calculation and give a realistic example of its application to a hypothetical German life insurer.

Introduction

According to the Article 122 of the Solvency II Directive [1], the insurance companies which use an internal model should calculate their Solvency Capital Requirement (SCR) via a full probability distribution forecast:

Where practicable, insurance and reinsurance undertakings shall derive the Solvency Capital Requirement directly from the probability distribution forecast generated by the internal model of those undertakings, using the Value-at-Risk measure set out in Article 101(3).

The Article 121 of the Solvency II Directive outlines the statistical quality standards which the calculation of the probability distribution forecast must comply with. In particular, the following requirements must be taken into account:

- *The methods used to calculate the probability distribution forecast shall be based on adequate, applicable and relevant actuarial and statistical techniques and shall be consistent with the methods used to calculate technical provisions.*
- *The methods used to calculate the probability distribution forecast shall be based upon current and credible information and realistic assumptions.*

- *Insurance and reinsurance undertakings shall update the data sets used in the calculation of the probability distribution forecast at least annually.*

In other words, Internal Models are not only expected to produce the probability distribution forecast, but required to produce it in an accurate, robust and auditable way.

In order to produce a probability distribution forecast, one would need thousands of economic balance sheets after one year. Each of those economic balance sheets would be a result of a valuation run starting after one year and consisting of thousands of valuation scenarios. Such a brute force nested simulation approach would solve the problem in principle. To be precise, it would consist of several thousands of outer real-world scenarios for projecting the company's assets and liabilities for one year and at least 1000 inner risk-neutral valuation scenarios at each of these outer nodes to revalue the assets and liabilities at year one. With the nested simulation approach being extremely demanding in terms of runtime, one is tempted to look for alternatives - so-called liability proxy modelling techniques - which deliver a probability distribution forecast within a more affordable simulation budget. These techniques include Replicating Portfolios [2], Curve Fitting [3] and Least Squares Monte Carlo (LSMC).

In this note, we show how LSMC yields a powerful tool for liability proxy modelling that can be used for the probability distribution forecast and displays the following outstanding properties:

- Accuracy of calculations
- Speed of calculations
- Consistent coverage of market, credit and insurance risks
- Robust and reliable validation
- Feasibility of process automation.

Acknowledgement: We express our gratitude to Adam Koursaris and Steven Morrison from Barrie + Hibbert Ltd. for a fruitful and rewarding collaboration on various aspects of the LSMC method.

1 LSMC Approach

The LSMC method, which dates back to Longstaff and Schwartz [4], is applied across a range of different applications from banking to the energy sector. In our paper, we discuss an application of LSMC to the calculation of the probability distribution forecast as required under Solvency II. In particular, we consider the task of calculating the probability distribution for the PVFP (Present Value of Future Profits) of a German life insurance company after one year.

Remark: Needless to say, the LSMC technique can be also applied to other economic balance sheet items such as the Best Estimate Liability (BEL) as well. Our LSMC application aims at the PVFP for two reasons. On one hand, the PVFP is of interest in its own right since it is the base item for MCEV valuations. On the other hand, key Solvency II items such as the SCR but also the Available Solvency Margin (ASM) can be derived by using the PVFP.

In the following subsections, we subsequently discuss the individual steps of the LSMC approach.

Step 1 - Identification and Choice of Risk Drivers

The PVFP displays a functional dependence on a very large number of economic (e.g. equity) and actuarial (e.g. lapse) state variables. In a first step we identify the risks that are relevant for your portfolio and choose an appropriate set of such economic/insurance state variables - so-called risk drivers - that explain these risk characteristics of your liabilities. An exemplary list of risk drivers could include:

- Short term interest rate
- Long term interest rate
- Interest rate volatility
- Equity value
- Equity volatility
- Corporate bond default rate
- Mortality Rate
- Lapse Rate
- Expenses

Step 2 - Expressing the functional relationship between PVFP and risk drivers

We are going to express the PVFP as a function of its risk drivers. In order to do so we use a large set of simulated data that contains information on the PVFPs that are obtained at various risk driver values.

- (i) We choose a large number n of possible risk driver values $r^{(j)} = (r_1^{(j)}, \dots, r_h^{(j)})$, $j = 1, \dots, n$, and assume that our model experiences the performance according to these risk driver values in the first year (we call each of those risk driver values an outer scenario). Then we evaluate the liabilities at each of those outer scenarios, but instead of choosing a full set of at least 1000 risk-neutral valuation scenarios (also called inner scenarios) we use just 10 inner scenarios emanating from each node. Thus this approach yields rough estimators $(y_j)_{j=1, \dots, n}$ for the corresponding node PVFPs (see Figure 1.1).
- (ii) We perform an ordinary least squares regression through the pairs $(r^{(j)}, y_j)_{j=1, \dots, n}$ in order to express the functional relationship between the risk drivers and the resulting PVFPs. As a result we get the relationship

$$f(r) = \sum_{j=1}^k \beta_j \varphi_j(r), \quad (1.1)$$

where the $\varphi_j : \mathbb{R}^h \rightarrow \mathbb{R}$, $j = 1, \dots, k$ form a set of k basis functions, that are able to cover the complex dynamics PVFP as a function of the risk drivers adequately (see Figure 1.2 for a one-dimensional example).

Even though each y_j would be a rough estimator for the corresponding exact liability value, the regression generates a refined and more accurate liability value where the single errors in the y_j are being balanced by the defining properties of the least squares estimator for $\beta = (\beta_1, \dots, \beta_k)$ (see Figure 1.2 for a one-dimensional example). This leads to a significant reduction of the number of scenarios and runtime required, as even 20.000 scenarios can lead to very promising results.

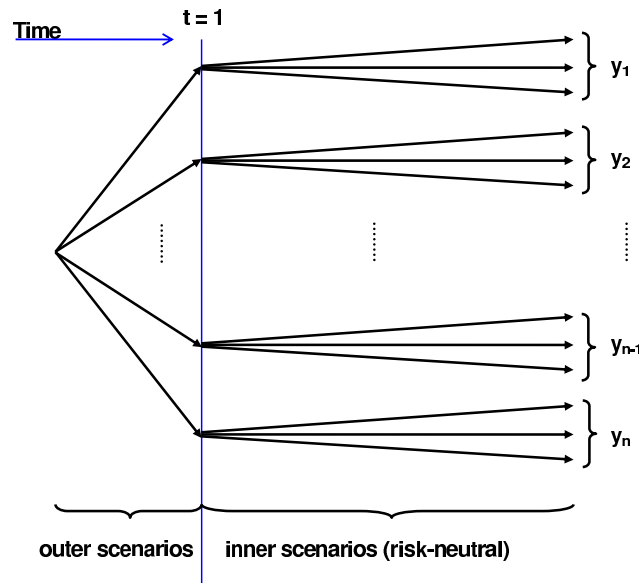


Figure 1.1: Split-up of scenarios in the LSMC approach.

In order to provide for a good quality of the liability function over a broad range we do not only use a large number of outer scenarios (≥ 5000) as input for the regression but furthermore choose their starting positions $r^{(j)} = (r_1^{(j)}, \dots, r_h^{(j)})$, $j = 1, \dots, n$, in a way that covers the space of possible risk driver values *evenly* instead of using traditional nested stochastic scenarios (differing only by the small number of inner scenarios employed) where the outer fitting scenarios would be the one year real world paths (as illustrated in Figure 1.3).

Step 3 - Calculation of Probability Distribution Forecast

Having expressed the functional relationship between the risk drivers as key determinants for the risk exposure in one year and the resulting PVFP values, we make use of the liability function for estimating the PVFP distribution in year one.

Therefore, we evaluate the polynomial function by using a set of, say, 100.000 one-year real world scenarios, i.e. 100.000 realizations of the risk drivers, featuring the desired dependencies between the relevant risk drivers. Finally we obtain the probability distribution forecast from the evaluation.

Step 4 - Validation of the Results

LSMC allows for explicit and implicit ways to validate the overall fitting process and the quality of the liability function, and hence the quality of the estimate of the PVFP distribution itself:

1. The most simple and direct way to validate the quality of the fitted function is to compare its point estimates for certain risk driver values with the corresponding PVFPs for these risk driver values obtained by a full Monte Carlo simulation with at least 1000 valuation scenarios.
2. We can use 2- and 3-dimensional plots that show the behaviour of the liability function in 1 or 2 risk drivers while all others being kept constant for validating the liability functions by judging and verifying the behaviour of the function in these risk drivers and their interplay.

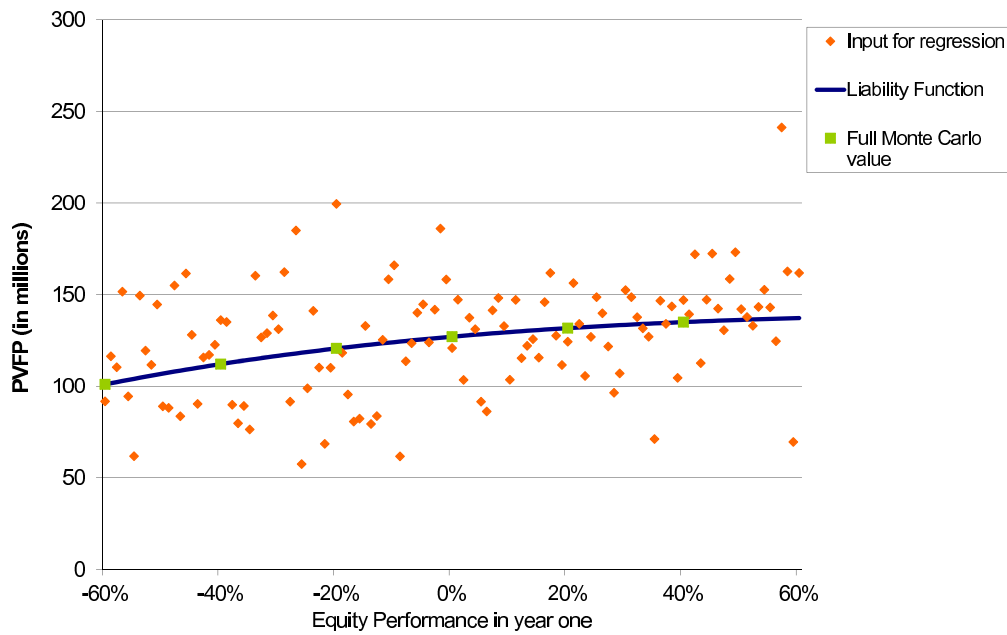


Figure 1.2: One-dimensional illustration of the LSMC approach, that plots the PVFP in dependency of the equity performance in year one.

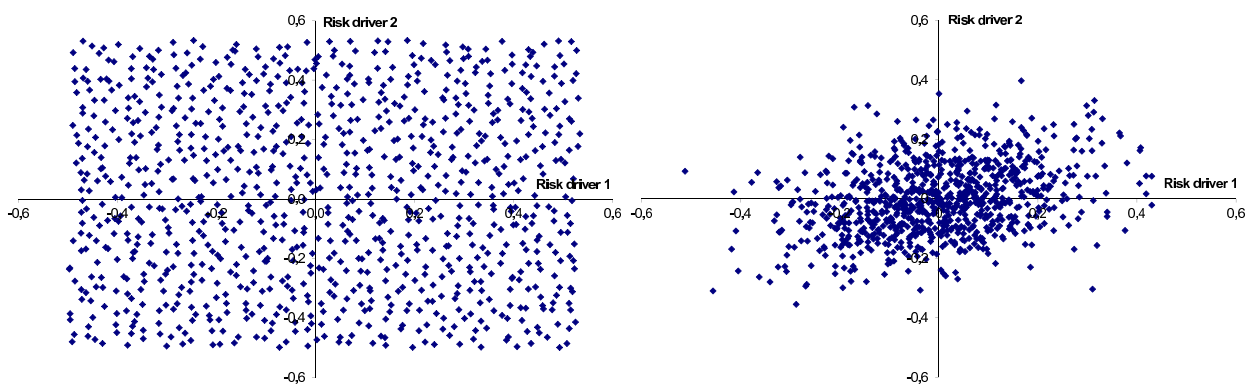


Figure 1.3: Two-dimensional example for evenly spread positions of risk drivers for fitting (left) and their real-world realizations for evaluation purposes (right).

3. There are simple but powerful statistical methods for LSMC - so-called Jackknife methods - that can even yield standard deviations and confidence intervals for the point estimates of the function and resulting values such as estimates of the VaR or Expected Shortfall. They can easily be calculated without re-running the full stochastic model and can be used to get a sense for the robustness and stability of the resulting estimates and detect errors in the overall fitting procedure.

2 LSMC Case Study

In this section we provide a detailed LSMC case study dealing with the LRA, a fictitious German life insurer, and subsequently applying the steps described above.

The LRA has been writing traditional annuity business since the 1950's and gradually diversified its portfolio over the years to include endowment, level-term and unit-linked business.

The balance sheet of the LRA at projection start is a typical one for a German life insurance company. In particular, its asset mix is as follows: 86% of the assets are invested in government and corporate bonds, 5% of the assets are invested in equities, 7% are invested in properties and 2% are invested in cash. The total fund accounting value of assets amounts to EUR 11 billions.

The liabilities include EUR 2 billions in unit-linked reserves and EUR 8 billions in conventional reserves, of which EUR 7.9 billions represent endowment and annuity contracts, whereas EUR 0.1 billions of reserves stem from level term contracts. The policyholder bonus reserve amounts to EUR 0.9 billions, whereas the statutory value of shareholder equity amounts to EUR 0.1 billions.

2.1 Calculation of Probability Distribution Forecast

We illustrate the LSMC technique by calculating the PVFP distribution of the LRA for year one and obtaining the company's SCR.

Step 1 - Identification and Choice of Risk Drivers

Based on the information given above, we identify the following risks as relevant for the LRA: Interest rate risk, interest rate volatility risk, equity risk, equity volatility risk, credit risk, lapse risk, mortality risk and longevity risk.

In order to define the market risk drivers, we refer to the corresponding parameters of the capital market model used: short rate level (serving as an indicator for the short term interest rate level), mean reversion level (serving as an indicator for the long term interest rate level), equity index level and equity volatility level. For the credit risk, we proceed analogously and refer to the credit factor driving the migration probabilities of corporate bonds in the credit risk model of LRA. For insurance risks, we consider multiplicative factors applicable to the individual insurance risks: lapse factor, mortality factor and longevity factor. For example, a mortality risk driver value of 0.5 represents 50% of the base mortality table.

Step 2 - Expressing the functional relationship between PVFP and risk drivers

We set ourselves a simulation budget of 50.000 scenarios. We spend this budget by running 5.000 outer real-world scenarios with 10 inner valuation scenarios each. There is a trade-off between the

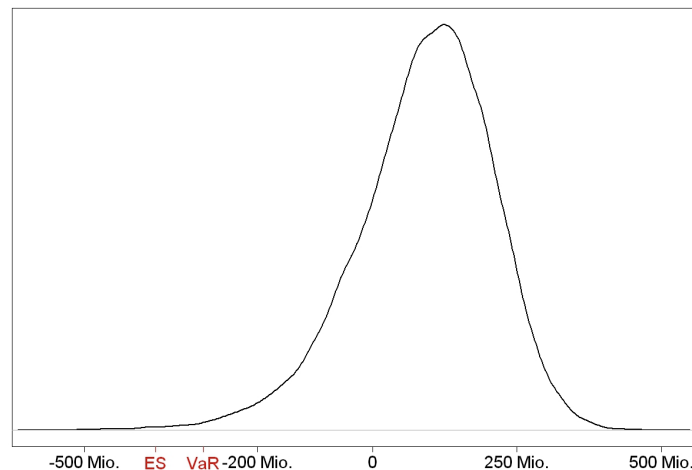


Figure 2.1: Year one PVFP distribution of the LRA including VaR and expected shortfall (both for 99,5% level).

aspiration to fill out the risk driver space as densely as possible and the aspiration to calculate the PVFP for each outer scenario as precisely as possible. Since the power of LSMC lies in its ability to fit the liability function across a range of possible future risk driver values, it is important to use a reasonably high number of outer scenarios.

Step 3 - Calculation of Probability Distribution Forecast

In order to get the PVFP distribution for year one we use outer real-world scenarios with a horizon of one year featuring a simple dependency structure for pragmatic reasons. In particular, we assume a zero correlation between market/credit risks and insurance risks. Furthermore, we use QIS5 correlations of insurance risks with one another.

Remark: The LSMC method would allow the risk manager to choose any dependency structure between the relevant risks, as long as the structure of choice were translated into a suitable package of, say, 100.000 real-world scenarios.

Evaluating our liability function from Step 2 on these real-world scenarios, we obtain a PVFP distribution of Figure 2.1. The 99.5% quantile of this distribution equals EUR -294 millions, whereas the expected shortfall amounts to EUR -366 millions. With the PVFP of LRA (at projection start) being equal to EUR 115 millions, we conclude that the SCR equals EUR 409 millions.

Step 4 - Validation of the Results

Figure 2.2 displays the results of the point estimates of the liability function and the corresponding exact values obtained by a full Monte Carlo simulation for a one-dimensional example that varies the interest rate mean reversion level only, all other risk drivers being constant at their unstressed levels. The corresponding 95% jackknife confidence level for these point estimators are plotted in Figure 2.2 as well.

Figure 2.3 contains further one-dimensional examples that illustrate the behavior of the liability

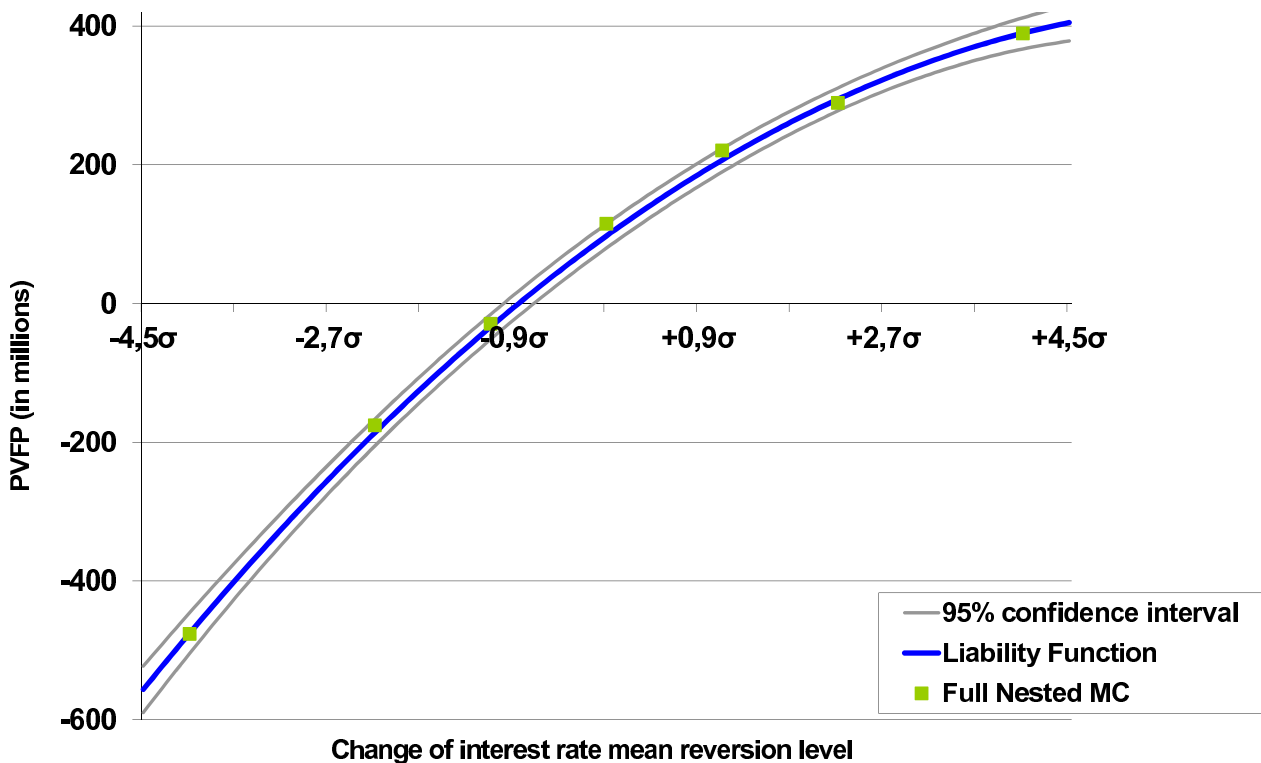


Figure 2.2: Point estimates from the liability function with corresponding 95% confidence interval and full nested Monte Carlo values.

function in different single risk drivers, all other risk drivers being constant at their unstressed levels. Such plots can be used to analyze whether the overall shape of the liability functions displays plausible dynamics. All plots display an overall decrease of the PVFP when the corresponding stress levels are increasing, which seems plausible, since an increase of volatilities, longevity prospects and corporate bond defaults should have a negative effect on the resulting PVFP. A particularly interesting effect occurs in the behavior of the PVFP in the interest rate volatility, where the PVFP at a very low level of volatility increases slightly with increasing volatility (and drops at higher volatility levels). We consider this behavior to be plausible by considering the fact, that - given a very low volatility and today's low interest rates - a moderate upward movement of the interest rate volatility level will lead to a rising PVFP since in such an economic environment higher volatilities will represent a prospect of improvement.

With the liability function at our disposal, we can analyze the interplay of different risks. We constrain ourselves to pairs of risk drivers since, say, a combination of 5 risk drivers would be neither easy to plot nor easy to grasp. Firstly, we plot the PVFP as function of the interest rate mean reversion level and the (historic) lapse. The lapse range covers stresses from +50% to -50% and the mean reversion level stresses range from $+4.5\sigma$ to -4.5σ , with σ being the standard deviation of the interest rate mean reversion process.

This surface (see Figure 2.4) can be explained via the following considerations:

- If both the interest rate mean-reversion level and the lapse rates are low, then the PVFP is extremely low, since the interest rate guarantees persist and the capital gains do not suffice.
- If the interest rate mean-reversion level is low, then rising lapse rates help reduce the guarantees and thus signify a rising PVFP.

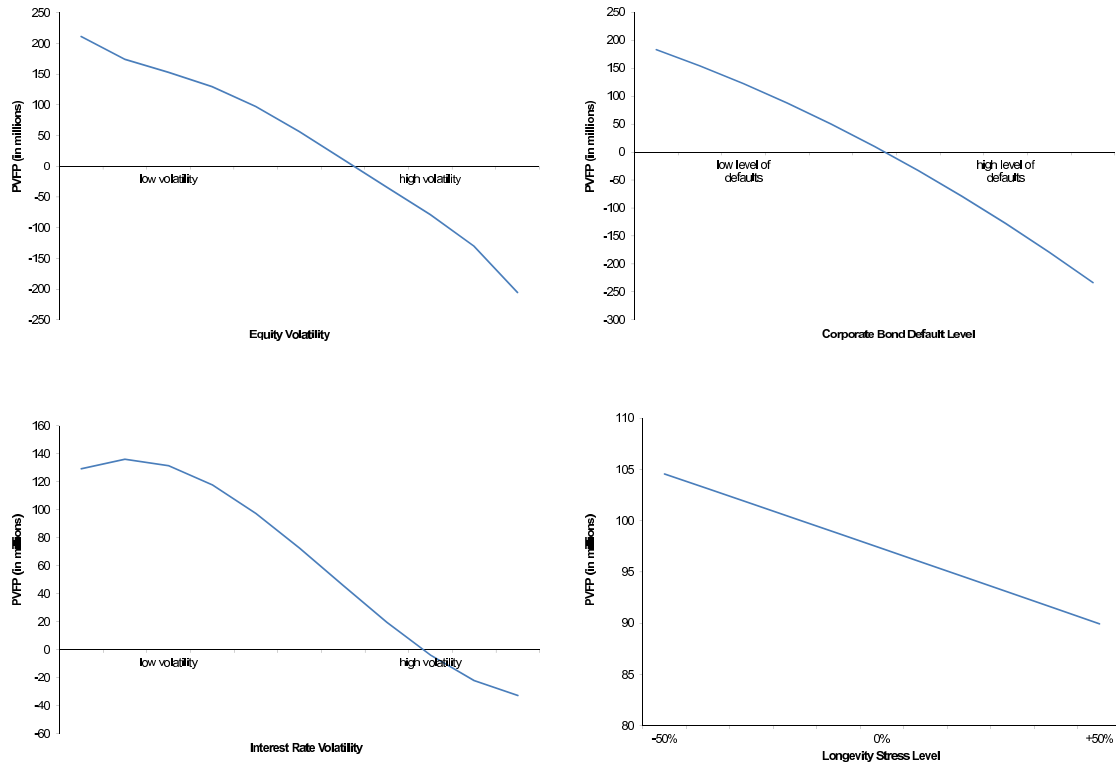


Figure 2.3: One-dimensional plots of the liability function that show the behavior of the PVFP in single risk drives: Equity volatility (top left), interest rate volatility (bottom left), corporate bond default level (top right) and longevity stress level (bottom right).

- At high lapse rates, a rise of the interest rate mean-reversion level leads to a significant PVFP increase due to rising capital gains and falling cost of guarantees.
- At a high interest rate mean-reversion level, decreasing lapse rates signify an increase of assets under management and thus lead to a rise of PVFP.

Summary: All three validation criteria - comparison of estimated PVFP against true Monte Carlo values, analysis of PVFP dynamics in 1 and 2 risk dimensions and calculation of confidence intervals for the liability function - indicate that LSMC produces results of high quality. The liability function can be accurately computed and gives no indications for errors in the fitting process. Thus, risk measures such as VaR or Expected Shortfall obtained via this method are reliable for risk management purposes. Furthermore, the liability function allows for an economically sound analysis of the impact of various risk drivers upon the results.

2.2 Refined Approach: PVFP Dependency on Management Rules

Using LSMC, we can not only analyze how the PVFP or the SCR depends on its risk drivers, but also explore the dependency of the PVFP or the SCR on the (parameters of) the management rules. In our example, we consider the equity backing ratio (EBR). At LRA, the EBR management rule is very simple - in fact, it keeps the EBR constant over all the projection years and scenarios. In other words, the only degree of freedom available is the EBR parameter.

We extend our list of risk drivers in order to cover the EBR parameter, so that this parameter is

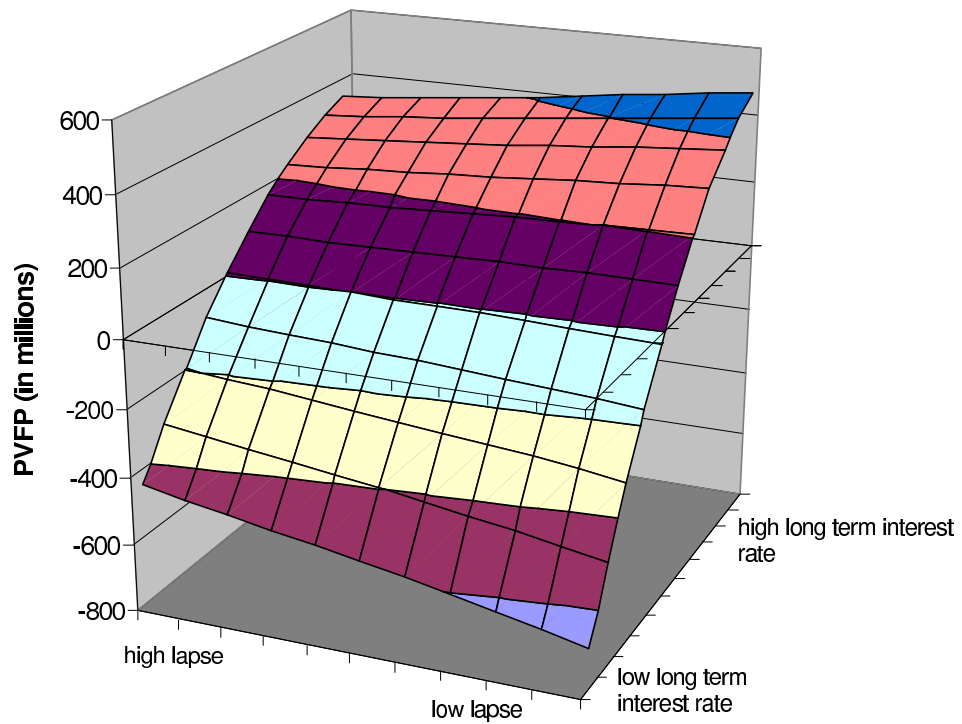


Figure 2.4: Two-dimensional risk dependency surface.

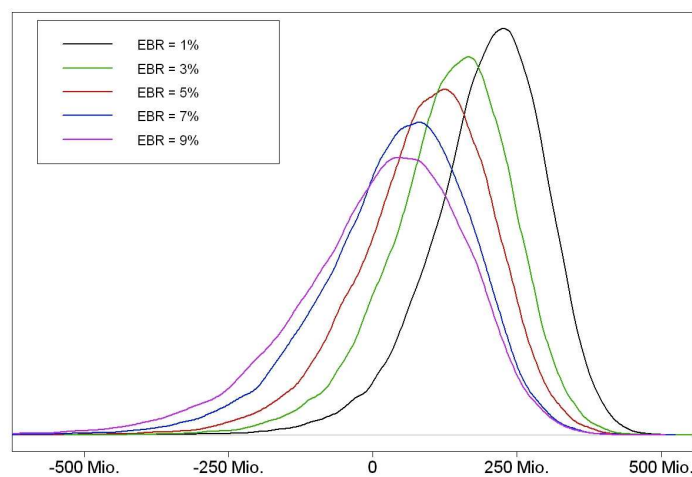


Figure 2.5: Year one PVFP distribution of the LRA for different equity backing ratios.

included in the outer scenarios and the regression and thus appears as an additional argument of the liability function. If we evaluate the function we can set this parameter to its base value while simulating the other risk drivers with their joint real world realisations as stated above. This leads to the distribution we discussed above. However, we can also use the information from the fitting across a range of different EBR values and change the value of the EBR parameter to another valid value. The result is again the full distribution of the PVFP - derived under the assumption of a different management rule parameter.

With the EBR = 5% in our previous example, we assume that the LRA management could set this parameter anywhere between 1% and 9%. Using LSMC, we obtain the following results:

Equity Backing Ratio	99, 5% VaR	SCR
1%	-143	258
3%	-222	337
5%	-294	409
7%	-391	506
9%	-482	597

Table 2.1: Distribution of PVFP (in millions) for different EBRs

These results can be explained in the following way: If a relatively high constant EBR is used, then it poses a substantial strain upon the shareholder in burn-through scenarios, whereas about 90% of the capital gains in good scenarios are allocated to policyholders.

The densities for each individual choice of the EBR parameter are visualized in Figure 2.5.

3 Comparison with other methods

We have seen that LSMC offers a powerful framework for calculating the probability distribution forecast in a very accurate and flexible way. In this section we summarize the outstanding features and benefits of the LSMC approach and compare it to Replicating Portfolios (RP) and Curve Fitting (CF):

- **Speed:** We found that an affordable simulation budget of 50.000 scenarios leads to LSMC results of high quality. The fitting can be carried out in less than 1 hour, the evaluation of the liability function only takes few minutes and the validation can be performed in less than 1 hour.
- **Simple choice of outer fitting scenarios:** Both RP and CF can only be calibrated against a rather small number of fitting scenarios. Hence choosing those scenarios involves expert judgement and can be quite cumbersome. In a contrast to this, LSMC does not require expert judgement, but thousands of fitting scenarios that are just evenly spread over the possible range of risk driver values in order to equip the liability function with as much information as possible.
- **Validation:** LSMC allows for the calculation of explicit confidence intervals of the liability function and the resulting capital requirements such as VaR and expected shortfall. Thus one can hence give a quantitatively precise answer to the question "How good is the overall quality of the VaR/expected shortfall estimate?"
- **Accuracy:** Using the least squares regression technique to cancel out the sampling errors

in the individual point estimates enables us to use a large quantity of outer fitting scenarios (≥ 5000) while RP and CF are both calibrated against only few outer fitting scenarios. The LSMC approach thus provides a lot more information on the dynamics and dependencies of the PVFP as a function of the risk drivers than RP and CF. This leads to a more thorough interpolation between those outer scenarios which is truly based on sampled data and does not involve "guessing" on how to interpolate between the sparse populated values of the outer scenarios.

- *Mathematical foundation:* There is a fundamental mathematical framework for LSMC that assures its theoretical convergence whereas both RP and CF have no such background since their quality depends heavily on the right choice of outer scenarios and assets/functions involved in the calibration/fitting process.
- *Flexibility:* Non-market risks and parameters of the management rules can be included in a canonical way. On the contrary, RP cannot necessarily deal with insurance risks such as lapse risk or mortality risk or even parameters of the management rules, since there are no liquid capital market instruments driven by the latter. Curve fitting will be theoretically able to deal with those kinds of risk drivers. However, due to the extremely small number of outer fitting scenarios it is more sensitive to the overall number of risk drivers and its quality will drop significantly in each additional risk driver being included in the model. In comparison, LSMC yields stable and robust results even for a large number of risk drivers and can deal with extension of the risk driver space with a relatively small additional amount of complexity in the liability function (this fact has already been reported by Longstaff and Schwartz).
- *Automation:* The overall LSMC process can be automated to a very high extent. All steps that are performed outside ALM-software (generating the outer fitting scenarios, fitting the liability function, calculating the probability distribution forecast and validation of the result) can be fully automated and need no expert judgement and human intervention.
- *Fees:* LSMC does not incur any significant software licence fees whereas RP one requires special software in order to construct replicating portfolios of high quality, which does incur significant licence fees.
- *Economic Insights:* By evaluating the liability function the interplay between different risk types can be analyzed and interpreted economically as described in Step 4. Such visualizations provide powerful tools to gain additional insight on the dynamics of the PVFP.

4 Outlook

The LSMC methodology is a powerful framework which can be used for a wide range of applications. We have shown how LSMC can be applied for the calculation of probability distribution forecasts over the 1-year-horizon, as required under Solvency II. We have seen in our case study that LSMC solves this task reliably and produces a number of additional insights which can be used for risk management purposes.

We would like to conclude this white paper by listing a number of related Solvency II applications of LSMC which we consider especially relevant in our work:

- Intra-year updates of the Solvency II calculation results,
- Medium-term economic capital planning over a 3-5y horizon,
- Group aggregation of SCR across life, health and non-life businesses,
- Development of management rules based on Solvency II coverage ratios.

We believe that especially the latter application is of fundamental interest. With Solvency II set to significantly influence business decisions for years to come, management rules must be able to take Solvency II into account, if they are to pass the Use Test.

Bibliography

- [1] Solvency II Directive 2009/138/EC.
- [2] J. Oechsli et. al., 2007. *Replicating Embedded Options*, Life & Pension Risk February 2007.
- [3] S. Carroll, C. Hursey, 2011, *Efficient Curve Fitting Techniques*, Life Conference and Exhibition.
- [4] F. Longstaff, E. Schwartz, 2001, *Valuing American options by simulation: A simple least-squares approach*, Rev. Financial Stud. **14**, pp. 113-147.

About Milliman

Milliman is among the world's largest independent actuarial and consulting firms. Founded in 1947, the company currently has 56 offices in key locations worldwide employing more than 2,500 people.

Milliman in Europe

Milliman maintains a strong and growing presence in Europe with 250 professional consultants serving clients from offices in Amsterdam, Brussels, Bucharest, Dublin, Dusseldorf, London, Madrid, Milan, Munich, Paris, Warsaw, and Zurich.

www.milliman.eu



Contact

For more information, please contact your local Milliman representative, or one of the authors:

Michael Leitschkis
michael.leitschkis@milliman.com
Tel. +49 211 88231 595
Fax. +49 211 88231 520

Mario Hörig
mario.hoerig@milliman.com
Tel. +49 211 88231 596
Fax. +49 211 88231 520

Milliman GmbH
Speditionstrasse 21
40221 Dusseldorf
Germany