Prepared by:

Mario Hörig
(Düsseldorf)

Karl Murray
(Dublin)

Eamonn Phelan
(Dublin)

Michael Leitschkis
(Düsseldorf)

Milliman

An application of Monte Carlo proxy techniques to variable annuity business: A case study

November 2013

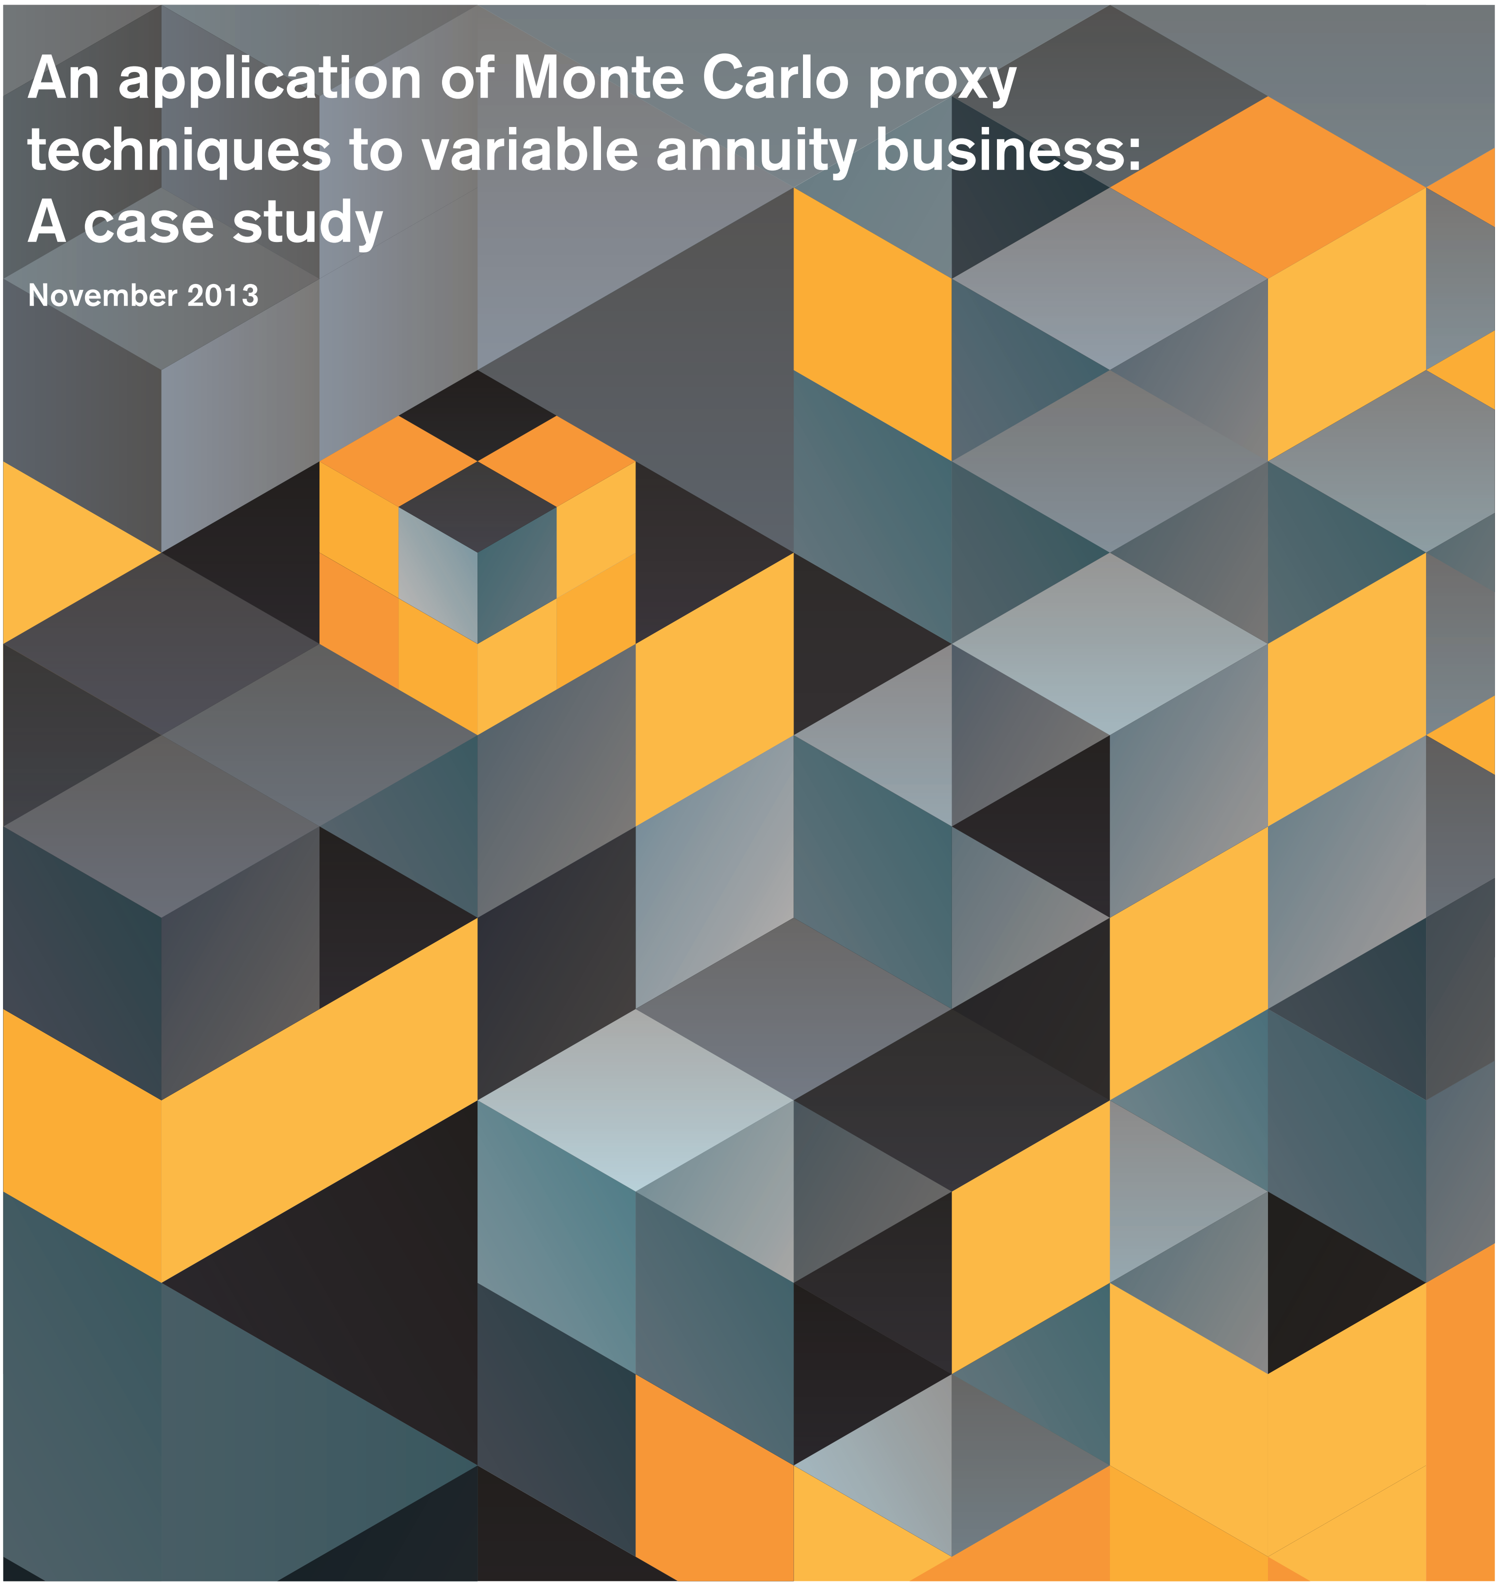

Milliman

TABLE OF CONTENTS

1 EXECUTIVE SUMMARY

1.1 BACKGROUND AND OBJECTIVES

Insurance and reinsurance undertakings face an ever-increasing demand from their various stakeholders for a deeper understanding of financial results and risk profiles. Meeting this demand essentially places huge pressures on existing financial modelling capabilities. Over time, the resource requirements of these systems (both in terms of cost and time) may become excessive.

A number of potential solutions to this problem have emerged. Amongst these is the suite of 'proxy modelling' techniques. The primary objective of this paper is to consider a case study application of one of these techniques to variable annuity (VA) business, namely the Least Squares Monte Carlo (LSMC) technique. We have chosen to base the case study on VA business as the complexity of this business generally requires very sophisticated and time-consuming modelling and is therefore an ideal candidate to illustrate the applicability of the LSMC proxy modelling technique. In demonstrating that the LSMC proxy modelling technique may be successfully applied to VA business, we also indirectly demonstrate that LSMC can handle the complexity associated with other types of business.

The case study covers the calculation of capital requirements as at the balance sheet date (and particularly in the context of Solvency II, although the conclusions have wider application to economic capital modelling in general). Consideration is also given to the extension of the LSMC method to other contexts such as Own Risk and Solvency Assessment (ORSA).

In this paper, for simplicity, we consider only the modelling of liabilities arising from VA products. We do not consider the modelling of the assets.

1.2 SCOPE AND STRUCTURE OF THIS PAPER

This paper begins with an introduction in Section 2 to the modelling problem faced by many insurers and reinsurers at present, together with its underlying causes.

Section 3 goes on to describe current modelling capabilities in more detail, particularly in the context of nested stochastic modelling techniques.

Proxy modelling techniques, and the reasons for choosing to apply such techniques, are covered in Section 4.

We highlight the LSMC technique in Section 5, providing details on how it may be applied, its theoretical background, and its key inputs, before going on to discuss the main advantages and limitations to bear in mind.

Section 6 provides background information on variable annuities as a product class together with a high-level description of the case study (data point and assumptions) underlying this paper.

In Section 7, we calibrate the LSMC proxy model to the case study data and assumptions and demonstrate how to validate the output from the model.

Projections that cover a multi-year time horizon generally form part of the modelling requirements of insurers and reinsurers. We discuss the application of the LSMC technique for such projections in Section 8.

We discuss in Section 9 how LSMC may be applied to capital regimes other than Solvency II, such as the existing Solvency I regime in Europe.

Section 10 contains numerous references which underlie many aspects of our research.

Full details of the case study and assumptions are contained in the appendix.

1.3 MAIN FINDINGS

This case study successfully demonstrates how the LSMC technique may be calibrated to model VA business, bearing in mind that in successfully applying this technique to such a complex product structure we have indirectly demonstrated that it will also work successfully for other product types.

LSMC is an objective and robust methodology for the determination of liability and capital values. In working through the calibration process it is possible to isolate instances of the need for expert judgment and to minimise reliance on such judgment.

Validation of the model calibration is critically important. We have found, through the application of confidence intervals, that it is possible to satisfactorily validate model outcomes under this proxy methodology.

Path dependency, dynamic policyholder behaviour, and dynamic management actions may add significant complexity to the modelling challenges of undertakings. Considering dynamic policyholder lapse behaviour as an example, we have shown how LSMC can take this into account without significant loss of accuracy.

LSMC may be used in order to address modelling challenges in a wide variety of real-life situations, from best estimate liability valuation to capital calculations, both at a given point in time as well as across multiple time periods (as illustrated by the case study in this paper). Hence, it is a methodology that has wide application as well as the necessary characteristics to address the modelling needs of a wide variety of undertakings.

2 INTRODUCTION

2.1 BACKGROUND

Insurance and reinsurance undertakings face an ever increasing demand from their various stakeholders for a deeper understanding of financial results and risk profiles. Innovation, coupled with consumer demand, has also driven complexity in product design, which has in turn led to the need for more sophisticated modelling techniques. Developments in financial reporting and capital standards have further added to this need.

In many cases the proper analysis of liabilities at a point in time requires the use of stochastic modelling techniques. Furthermore, understanding the future evolution of an insurer's liability profile can require the use of nested stochastic modelling techniques, whereby at each node of a projection a further nested stochastic projection is used, as outlined in Section 3 below (and illustrated in Figure 1). Such techniques tend to require significant resources if acceptable run times are to be achieved.

FIGURE 1: NESTED STOCHASTIC SCENARIOS

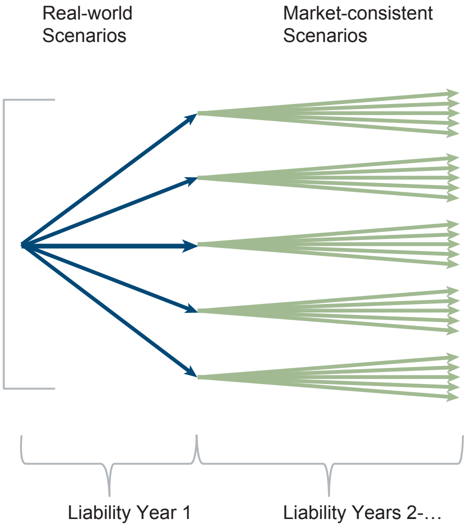

In order to continually manage model run times and to keep them within what might be considered to be practical time frames, a number of approaches have been adopted, including:

- Reliance on continued development of ever more efficient hardware and software
- A larger volume of hardware, e.g., use of grid computing and cloud computing
- Reductions in the numbers of scenarios and stresses examined
- Use of representative liability model points rather than the entire portfolio
- Simplification of cash-flow models
- Use of proxy modelling techniques

These approaches may be adopted in isolation or in combination with one another in order to achieve the desired outcome. It is the use of proxy modelling techniques that will be the focus of this paper, in particular the LSMC technique.

2.2 OBJECTIVES

The primary objective of this paper is to consider a case study application of the LSMC technique to VA business. LSMC, which is explained in further detail in Section 4.2.1 and Section 5 below, is a flexible method that can be applied in a range of different situations. However, VA business in particular exhibits characteristics that are very well suited to the application of LSMC. For example, the typical investment guarantees offered with such products will most likely necessitate a stochastic approach when assessing risk exposures and capital requirements. Furthermore, frequent up-to-date information is required as VA liabilities can change rapidly over time. With the growth of VAs in the European market in recent years, following significant growth in the past in other markets (such as the United States and Japan), writers and reinsurers of such business are facing significant model resource demands.

This case study analysis considers two typical VA product structures and applies LSMC to the calculation of the Solvency II Solvency Capital Requirement (SCR). For simplicity, representative model point policies have been constructed for these two product designs. The case study provides a step-by-step guide to the implementation of the LSMC technique as well as benchmarking the outcomes for accuracy against a fully stochastic calculation in order to demonstrate the accuracy and robustness of the proxy modelling technique. For the purposes of this case study, it is primarily market risk that is modelled, although lapse risk is also modelled to the extent that exposure to lapse risk is impacted by the performance of the underlying funds relative to the guaranteed benefits and market conditions in general.

While the majority of the case study considers capital at the balance sheet date and in the context of Solvency II, consideration is also given to the extension of the LSMC method to other contexts such as the ORSA (see Section 8.1.1) and Solvency I (see Section 9). Although Solvency I is a European regulatory framework, the conclusion of this paper should be equally applicable to other jurisdictions.

In this paper, for simplicity, we consider only the modelling of liabilities arising from VA products. We do not consider the modelling of the assets. However, LSMC may also be applied to the asset side of the balance sheet (for example, the valuation of American options). Of course, it is also possible to apply LSMC as a proxy for assets minus liabilities (i.e., the ‘surplus’ or ‘net asset value’) instead of attempting to create separate proxy valuations for the assets and liabilities. However, in general, insurers and reinsurers tend to model these separately. In addition, both assets and liabilities are individually very large numbers, whereas surplus is much smaller and generally much less stable.

3 NESTED STOCHASTIC MODELS

3.1 THE CASE FOR NESTED STOCHASTIC PROJECTIONS

While point-in-time liability valuations generally require nothing more complex than a normal stochastic calculation, nested stochastic models are particularly useful when it comes to projecting either financial results or capital requirements over time. This challenge may arise under a number of different guises, such as:

- Solvency II
- IFRS
- C3 Phase II and Phase III capital requirements
- Actuarial Guideline 43 (AG43)
- Market-consistent embedded value
- Economic capital

Essentially, whenever an insurer or reinsurer needs to produce a distribution of projected outcomes over time, where reserves or capital need to be determined using stochastic techniques, there will be a strong case for nested stochastic modelling. Intensive modelling is a characteristic of financial projections for guaranteed business and, in particular, VA business.

3.2 NESTED STOCHASTIC MODELLING AND SOLVENCY II

Nested stochastic techniques will not be needed for point-in-time calculations under the Solvency II standard formula. Instead, the need for nested stochastic simulations is likely to be primarily an issue for internal model or partial internal model users. It may also prove useful for the purposes of the ORSA to project future capital requirements (whether using the standard formula or another modelling approach). Internal models will need to project a distribution of expected future capital requirements. In the case of VA business, liability and capital calculations generally necessitate stochastic techniques. Therefore, the development of each point of the capital requirement distribution is likely to need a further (nested) stochastic calculation.

In other words, 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 market conditions over a one-year time horizon and, in the case of Solvency II, 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.

For standard formula purposes, in order to calculate their current capital requirement, it is sufficient to apply a single stress factor to individual risks, so there is, in a sense, just a single outer scenario. However, under the ORSA requirements, there will be a need to project capital requirements into the future under a range of potential outcomes. In a sense, this is similar to a nested stochastic approach, depending on the range of future outcomes to be projected.

The production of full-scale nested stochastic projections poses huge challenges in terms of run time, as there are so many individual stochastic paths required to be projected. For example, an insurer or reinsurer using an internal model under Solvency II may run something on the order of 100,000 outer scenarios and 5,000 inner scenarios. This would necessitate 500,000,000 stochastic paths to be projected with a one-year period projection horizon in the outer scenarios.

In the context of Solvency II, the LSMC proxy modelling technique may be applied to individual risk modules under a partial internal model approach and later aggregated with other risk modules calculated under a standard formula approach or indeed other risk modules calculated in other ways such as alternative proxy modelling approaches or full nested stochastic projections. Before discussing this in detail, it is useful to briefly outline some alternative proxy modelling techniques, as set out in the next section.

4 PROXY MODELLING

4.1 THE CASE FOR PROXY MODELS

Life insurers and reinsurers offering investment guarantees will typically monitor their exposures to market risk, in particular, on a daily basis. The value of the guarantees they underwrite can change rapidly as investment markets move. Companies will make use of thousands of economic scenarios to assess their liability exposures. In addition, many writers and reinsurers of guaranteed business will hedge their economic exposures using derivative instruments. The combination of needing to model relatively complex assets and liabilities and requiring regular updates of their risk exposures poses a practical problem.

Such companies need a fully flexible valuation process allowing for timely and frequent delivery of results to management. These deliverables generally comprise key liability values as well as solvency and economic capital figures and include various sensitivities of the results to market changes such as equity market falls or interest rate changes.

The valuation of complex insurance liabilities such as VA products generally requires stochastic scenarios in order to determine the market value of the embedded options and guarantees. Such valuations typically require huge numbers of simulations in order to properly assess the market value of these products, particularly if there is an element of path-dependence in the guarantee.

For example, if determining the SCR in the context of Solvency II, (re)insurers using internal models face the challenge of evaluating their liability portfolios for thousands of real-world market outcomes over a one-year time horizon. This typically results in the need for nested stochastic simulations running to millions of individual scenario projections.

FIGURE 2: NESTED STOCHASTIC SCENARIOS

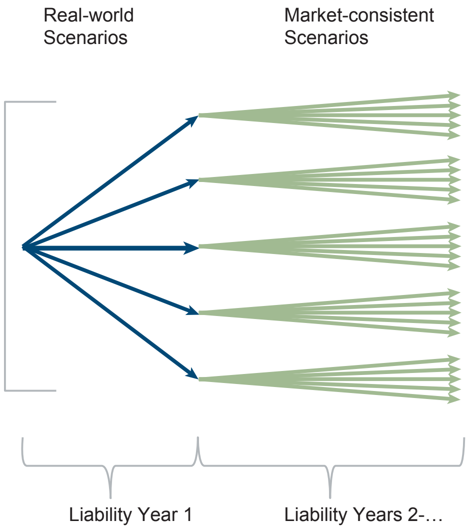

This approach can lead to real challenges for companies in managing run times and in having the IT capability to perform the modelling. These challenges cannot necessarily be overcome through additional hardware. Hence, various proxy modelling techniques have become increasingly popular.

Proxy models tackle this problem in a number of ways–for example, by providing a functional relationship between the economic variable under consideration, e.g., the best estimate liability (BEL), and the market conditions at time t=1. Once such a relationship has been determined, the valuation of the liabilities becomes a much less onerous task as it can be performed simply by evaluating the proxy. Hence, deriving a real-world view on the distribution of the liability value basically means evaluating the proxy for a large set of real-world realisations of these market conditions.

Deriving reliable and robust proxy models which require only minimal expert judgment is a major challenge facing (re)insurers. Actuarial projection models incorporate the projection of complex liability products such as VAs, including complex dynamic management actions and policyholder behaviour, and a successful proxy model has to incorporate these effects and interactions. The following section outlines a number of alternative approaches.

4.2 PROXY MODELLING APPROACHES

Over recent years, a number of proxy modelling approaches have begun to be widely used by life insurers and reinsurers. These include the following techniques:

- Least Squares Monte Carlo
- Curve fitting
- Replicating portfolios

This research paper investigates the use of LSMC in the context of calculating and projecting capital requirements for VA business. However, before getting into the detail of LSMC and its application to VA business, the following sections provide a brief overview of each of the techniques mentioned above.

4.2.1 LEAST SQUARES MONTE CARLO

The LSMC method is applied across a range of different applications, from banking to the energy sector. It is a standard numerical method for option pricing where the option is exercisable at points prior to ultimate maturity, otherwise known as an American-style option.

The accuracy of results obtained using the LSMC proxy approach (relative to using a full stochastic approach) can be verified in a practically robust and statistically sound way. The LSMC method generally requires much less manual intervention than some of the alternative proxy methods and may also give valuable insights about the interplay of different risk drivers.

In general, LSMC has the following beneficial properties:

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

LSMC can be used in any situation where the complexity of liabilities necessitates a stochastic valuation approach. For example, where there are embedded options in contracts or scope for future management actions, in addition to the modelling of investment guarantees.

Further details on the implementation of the LSMC approach are provided in Section 5.

4.2.2 CURVE FITTING

Whereas LSMC considers fitting a regression model to a large number of inaccurate valuations to remove the statistical sampling error, curve fitting considers a small number of very accurate calculations. At a high level, the curve-fitting process develops as follows (where, in this example, the objective is to project a BEL):

- Recalculate the BEL under a small number of instantaneous shocks. This involves running a large number of inner scenarios for each outer scenario and averaging the results of the inner scenarios at each outer scenario node point.
- Fit a multidimensional polynomial to the BEL values by regression techniques. The fitting itself requires a predetermined view about the structure of the polynomial that is used for this purpose.
- The validation of the curve fitting can only be carried out by performing out-of-sample tests where the true BEL value for an arbitrary outer scenario is derived by using a large number of inner scenarios and comparing this with the estimate produced by the curve fitting function.

The method can work well with a small number of risk drivers. The main differences relative to LSMC are:

- The calculations can quickly become quite onerous because fitting a polynomial with n roots requires the use of valuations for up to n different outer scenarios. In circumstances where there are complex interactions and many risk drivers, this will lead to a large number of valuation scenarios involved.
- The use of a predetermined polynomial could introduce an intrinsic error if the polynomial structure is unsuitable, e.g., if it misses certain terms. The choice of the right polynomial involves in-depth knowledge of the effect of the single risk drivers and their interactions for the product under consideration.
- Choosing the right fitting points is crucial for the quality of the resulting fit and it can be difficult to choose these fitting points appropriately.
- The validation and assessment of the error of the fit is limited to brute-force validations that involve large numbers of scenarios.

It can be shown that LSMC and curve fitting essentially converge to the same answer under certain conditions. In addition, LSMC can be used for some risk drivers and curve fitting for others and the two results combined. There are, of course, limitations with this latter approach (for example, it generally fails to capture interactions of the risk drivers between those modelled using LSMC and those modelled using curve fitting), but it can still work well in certain circumstances.

4.2.3 REPLICATING PORTFOLIOS

The basic idea of replicating portfolio techniques is to replace the complexity of the liability calculation with a proxy portfolio of market instruments whose value closely replicates the change in the value of the (re)insurance liabilities across a range of potential outcomes. The valuation of the market instruments could be much more straightforward than the valuation of the liabilities directly, in which case the nested stochastic problem may be very much simplified.

The general process involved in finding a replicating portfolio is as follows:

- Choose a set of potential assets
- Run a range of scenarios through the liability cash-flow model
- Solve for the weights in each of the assets that would replicate the liability cash flows
- Revalue the asset portfolio under the desired set of scenarios as a proxy for the liability valuation

Replicating portfolios can be useful for dealing with market risks and, in particular, can have immediate applications for ALM exercises and hedge asset selection. However, it can be more difficult to use them in relation to non-market risks as the assets won't necessarily be sensitive to the types of non-market risk factors that drive the liabilities. It can be quite challenging to choose a good replicating portfolio, depending on the complexity of the liabilities under consideration as well as factors such as the availability of suitable assets in the market place, the liquidity of such instruments and the level of transparency in pricing.

4.2.4 GENERAL COMMENTS ON PROXY MODELS

Each of the approaches outlined above has its own unique features and may be more or less applicable depending on individual situations. One observation that is often made regarding both curve fitting and replicating portfolios is that neither of these techniques is straightforward, and when there are complex interactions between assets and liabilities it may be difficult to get an adequate level of fit. In addition, a significant amount of professional judgment needs to be applied under these two methods. Therefore, a lot of attention has recently focused on LSMC in an effort to address these issues.

Arguably, other methodologies may also be classified as proxy approaches, such as statistical methods which pre-select the 99.5% value-at-risk (VaR) outcome without running a whole distribution of stochastic cash-flow projections.

However, the examples discussed above may be considered amongst the most commonly used.

5 APPLICATION OF THE LSMC TECHNIQUE

5.1 HIGH-LEVEL INTRODUCTION TO THE TECHNIQUE

LSMC aims to determine a functional relationship between key parameters used for the evaluation of the liability. These key parameters are referred to as 'risk drivers,' and the corresponding liability value for the purposes of this case study is the BEL but could be any economic measurement of the product cash flows.

The overall LSMC process can be briefly summarised as follows:

Step 1: Evaluate the liabilities for a number of different deterministic joint positions of the risk drivers at time t=1 (the so-called outer scenarios) by using only a small number of valuation scenarios (so-called inner scenarios) per outer scenario. Typically this 'small number' is between 1 and 10 scenarios. The liability value per outer scenario is the mean of the liability values that have been simulated using the corresponding inner scenarios. Hence any liability value itself is a very imprecise estimate for the corresponding outer position.

This step basically collects a lot of (inaccurate) information concerning the relationship between risk drivers and liability values.

FIGURE 3: LSMC SCENARIOS

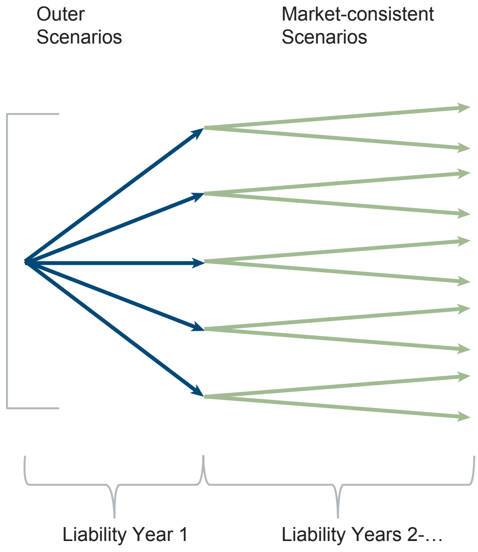

Step 2: Perform least squares regression based on the output from Step 1 to determine a functional relationship between risk drivers and liability values.

This step removes the remaining variance that has been left from Step 1 and uncovers the underlying relationship between risk drivers and liability values.

FIGURE 4: INACCURATE LIABILITY VALUES AND RESULTING LIABILITY VALUE CURVE

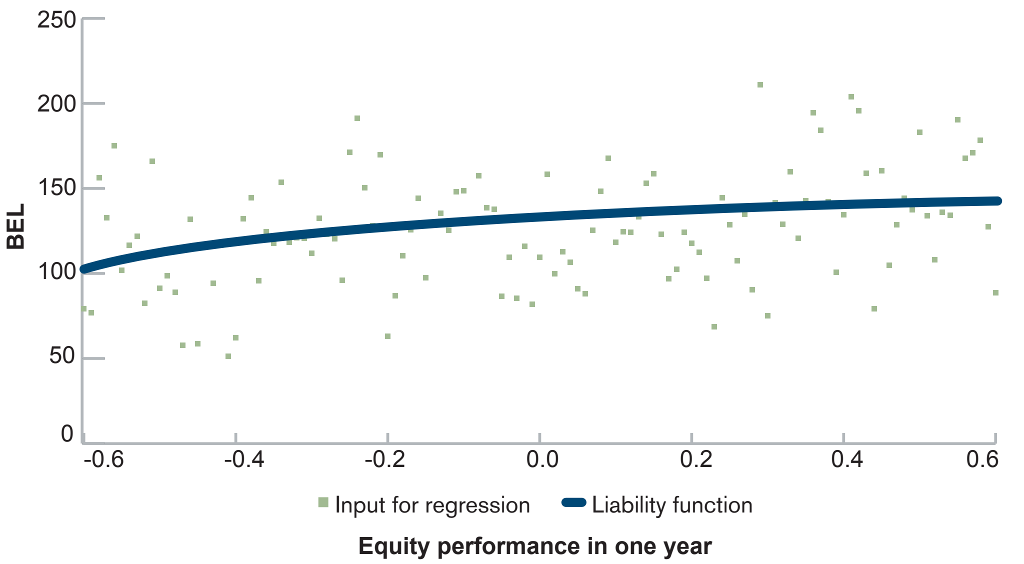

The overall aim is to significantly reduce the number of simulations involved by first evaluating the liabilities using only a limited number of inner scenarios, with the high residual level of 'noise' being removed in a second step by the least squares regression technique.

5.2 THEORETICAL BACKGROUND ON LSMC

5.2.1 LINEAR REGRESSION MODELS

LSMC is based on linear regression models. The fundamental idea of linear regression models is to model the relationship between a variable Y (the so-called *response variable* or *dependent variable*) and some explanatory variables $X_1,\ldots,X_m$ (so-called *independent variables*) via the relationship

$$Y = \beta_0 + \sum_{j=1}^{k} \beta_j \varphi_j(X_1,\ldots,X_m) + \varepsilon =: f(X_1,\ldots,X_m) + \varepsilon$$

where $\varphi_j : R^m \to R$, j=1,…,k, form a set of *k basis functions* and ε is the error variable (for simplicity we denote the overall function by *f*) . The model is called linear because the relationship between *Y* and the basis functions φ_j is linear. Note that the functions φ_j themselves can be non-linear, e.g., φ_1 $(X_1,\ldots,X_3)= X_1^2 * X_2^3$.

In our case, the dependent variable *Y* refers to the liability value while the independent variables $X_1,\ldots,X_m$ are the risk drivers parameterising the conditions at a relevant future time horizon (e.g., t=1 for a 1-year-VaR problem).

5.2.2 HOW TO CALIBRATE A LINEAR MODEL TO A SET OF OBSERVATIONS

Consider that we have made n different observations ($Y^{(i)}, X_1^{(i)}, \ldots, X_m^{(i)}$), i=1,…,*n* of the dependent variable and the corresponding independent variables, and want to calibrate a model represented by some basis functions φ_j , j=1,…,k to this data. This means that we aim at finding a set of coefficients $\beta_0, \ldots, \beta_k$ such that the resulting linear model

$$f := \beta_0 + \sum_{j=1}^{k} \beta_j \varphi_j$$

yields an optimal proxy for the observed data. Mathematically, this means that we aim at deriving a vector of coefficients ($\beta_1, \ldots, \beta_k$) that minimises the sum of squared residuals (RSS)

$$RSS = \sum_{i=1}^{n} \left(f(X_1^{(i)}, \ldots, X_m^{(i)}) - Y^{(i)} \right)^2$$

The solution $\beta = (\beta_0, \ldots, \beta_k)$ for this optimisation problem is the corresponding least squares estimator and it can be shown analytically that it has the following form:

$$\beta = \left(X^T X \right)^{-1} X^T y$$

where $X = \begin{pmatrix} 1 & \varphi_1\left(X_1^{(1)}, \ldots, X_m^{(1)}\right) & \cdots & \varphi_k\left(X_1^{(1)}, \ldots, X_m^{(1)}\right) \\ \vdots & \vdots & \ddots & \vdots \\ 1 & \varphi_1\left(X_1^{(n)}, \ldots, X_m^{(n)}\right) & \cdots & \varphi_k\left(X_1^{(n)}, \ldots, X_m^{(n)}\right) \end{pmatrix}$ and $y = \begin{pmatrix} Y^{(1)} \\ \vdots \\ Y^{(n)} \end{pmatrix}$.

The Gauss-Markov theorem states the most important properties of the least squares estimator and is of particular importance for the convergence of LSMC. According to the Gauss-Markov theorem, assuming that the errors ε are normal and independent identically distributed with mean 0, the estimator β as defined above is the Best Linear Unbiased Estimator (BLUE), i.e.,

- It is a maximum likelihood estimator
- It is unbiased
- It has the lowest variance among all possible linear and unbiased estimators

By virtue of the Gauss-Markov theorem, we see that, based on the underlying assumptions, the least squares estimator is the best estimate.

5.2.3 WHY LSMC WORKS

The LSMC approach derives a functional relationship between the risk drivers $X_1,\ldots,X_m$ and the corresponding liability value Y. First we derive rough estimates $Y^{(1)},\ldots,Y^{(n)}$ of the liability value for n different risk driver positions $(X_1^{(i)},\ldots,X_m^{(i)})$, i=1,…,$n$. Therefore, in the context of linear regression models, we can assume $Y^{(1)},\ldots,Y^{(n)}$ are only rough estimates because their values have been derived by using only a small number of valuation scenarios (e.g., each one might be the BEL calculated using only 10 risk-neutral inner scenarios instead of 1,000 or 2,000 such scenarios).

Using fewer valuation scenarios (or inner scenarios) per risk driver positions (or outer scenarios) allows for the generation of large numbers of inaccurate (or noisy) liability values within quite reasonable run-times. In accordance with the Central Limit Theorem, the error of each single estimate $Y^{(i)}$, i=1,..,n, i.e., the difference between $Y^{(i)}$ and the corresponding actual liability value (as derived by using, for example, 2,000 simulations) is asymptotically normally distributed with mean 0. Therefore, we are now in the context of linear regression models where we assume that there is a functional relationship between certain dependent variables (risk drivers) and an independent variable (liability value) with some additional noise (ε). Hence, by using a least squares regression with a proper set of basis functions the LSMC approach uncovers the true relationship between Y and the risk drivers $X_1,\ldots,X_m$ from the noisy input data. Basically, LSMC removes the errors ε and what remains is the true functional relationship between the liability value and risk drivers. For a detailed proof of the convergence of LSMC see Kalberer in the References section.

The least squares approach removes the noise from the input data and uncovers the true functional relationship.

5.2.4 EFFICIENCY OF THE TECHNIQUE

While we have seen that LSMC works in the sense that its estimates converge to the true values, it is worth mentioning why the approach of using many noisy liability values for the estimation of the functional relationship is particularly efficient. The reason for this lies in the square root convergence of the mean values themselves.

When deriving the actual BEL by using L valuation scenarios, we know from the Central Limit Theorem that the estimates based on L scenarios converge for $L \rightarrow \infty$ against the actual value at the speed of $1/\sqrt{L}$. Therefore, the overall simulation procedure ultimately becomes quite inefficient as more and more additional scenarios are required to get final convergence towards the actual value (see Figure 5, showing the convergence of a sample liability value over the number of valuation scenarios involved). In this example, while the actual value is 350, we have come quite close to this using very few scenarios (e.g., estimated liability value with first 10 scenarios is 386). While a small number of scenarios can generate a good estimate, it takes a significant number of scenarios to determine the actual value.

FIGURE 5: CONVERGENCE OF LIABILITY VALUES

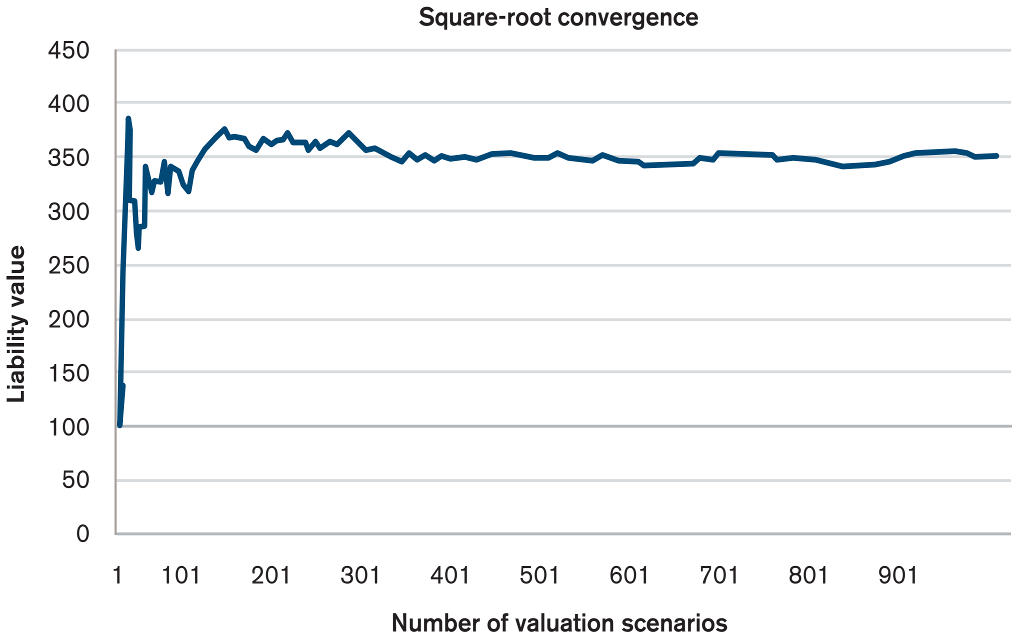

In summary, the LSMC approach uses only a small number of valuation scenarios and therefore produces noisy estimates of the liability data. The remaining noise is removed by the least squares regression. By using only a few inner scenarios per outer one we already have quite reasonable estimates of the corresponding liability values due to the square root convergence (see above). This allows us to use many outer scenarios for a given overall number of scenarios. Therefore, we use the scenarios involved particularly efficiently. Using many inner scenarios but just a small number of outer ones would be less efficient because most of the inner scenarios are used for slight improvements of the liability value estimates (where the remaining noise will be removed by the least squares regression anyway).

5.2.5 KEY INPUTS

The most crucial ingredient for a successful LSMC application is a proper selection of the basis functions that are used to derive the functional relationship between liability value and risk drivers. Note that, in our approach, the basis functions are special combinations of ordinary polynomials.

Examples:

- If the liability value displays a non-linear behaviour in a particular risk driver, we have to make sure that the functional relationship takes this into account by having cubic, quadratic or higher-order polynomial terms as part of the function.
- If there are interactions between certain risk drivers, we have to make sure to include proper polynomial cross terms in the function (e.g., $x_1 * x_2^2$ if there is such an interaction between risk drivers x_1 and x_2).
- Using too many and unnecessary polynomial terms will lead to *overfitting*, i.e., explaining effects that are not significant.

This leads to the following questions:

- How can we make sure to include all significant polynomial terms that are required and exclude insignificant ones?
- Given several functions $f_1, f_2, \ldots$ that could be used for the underlying linear model, which one is best for the given data?

In order to avoid any pitfalls related to these questions, we make use of well-known statistical procedures (called *model selection algorithms*) in order to detect the polynomial that best explains the effects that are shown in the fitting data (i.e., the rough liability values $Y^{(1)},\ldots,Y^{(n)}$ and the corresponding risk drivers $(X_1^{(i)},\ldots,X_m^{(i)})$, i=1,…,n), based on a well-established statistical criterion known as the *Akaike Information Criterion* (AIC). This criterion enables us to judge and compare the explanatory power of different liability functions by comparing their respective AIC values. The AIC value of a linear model f is defined as $AIC(f) := n * log(RSS(f) / n) + 2 * (k+1)$, where $k+1$ is the number of terms of f. Its judgment of the explanatory power of a certain polynomial for a set of calibration data considers the RSS of the model (as defined in Section 5.2.2.), i.e., it prefers models with small RSS because this is an indicator for a good fit, but simultaneously penalises the use of additional terms by adding 2*(k+1) with k+1 being the number of terms involved. The polynomial with the lowest AIC value is best for the given data because it minimises the RSS under the constraint to use as few polynomial terms as possible.

To select the best possible function for the relationship between the liability value and risk drivers we start with a trivial model (i.e., only the constant term) and iteratively add further terms. In each step of this selection procedure, we take the previous model (i.e., those polynomial terms selected in the previous steps), add different possible further terms, compare their AIC values, and choose the one with smallest AIC value as the new model. This procedure terminates after a number of steps as the AIC criterion penalises additional terms and hence will indicate that adding any further term won't be favourable after a certain number of polynomial terms have already been chosen.

Figure 6 illustrates this process by showing the following three stylised steps:

- Choice of a constant term at the start of the LSMC calibration process
- Addition of a linear term
- Addition of a quadratic term

FIGURE 6: MODEL SELECTION-EQUITY PERFORMANCE

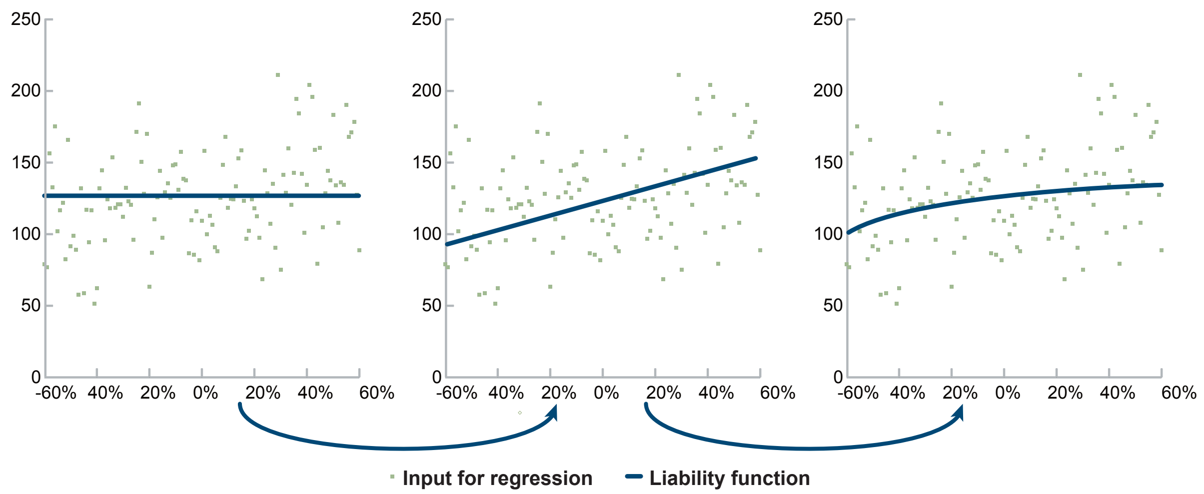

5.2.6 CONFIDENCE INTERVALS FOR LIABILITY ESTIMATES

LSMC not only allows for the efficient and robust estimation of a liability function but can also derive confidence intervals for its liability value estimates.

LSMC doesn't just answer the question, 'What is the liability value for a given set of risk drivers?' but also indicates how much confidence to have in these estimates.

Suppose we have derived a liability function *f*. For a certain risk driver combination $X_1,\ldots,X_m$ we get an estimate for the corresponding liability value by evaluating $f(X_1,\ldots,X_m)$ (i.e., a point estimate). In order to derive confidence intervals for the point estimate $f(X_1,\ldots,X_m)$ we apply a certain resampling procedure known as *jackknifing*. The jackknifing involves resampling of the liability function from the simulation data used for the LSMC fitting. Each resampling results in a new liability function and a new point estimate. The collection of new point estimates is used to derive standard deviations and confidence intervals for the point estimates of the liability function.

Further detail on this procedure is contained in Appendix A.3. The case study that is described below illustrates the use of confidence intervals for the judgment of the LSMC results.

5.3 LSMC PROCESS

The overall LSMC process begins by identifying the underlying risk drivers in order to make sure that all relevant risks are covered in an efficient way. Based on the risk drivers and their particular parameterisations, we create:

- *n* outer scenarios, i.e., joint positions of the risk drivers at the relevant future horizon (e.g., time t=1 if we are dealing with a 1-year-VaR problem)
- *N* inner scenarios per outer scenario, in order to evaluate the liabilities for the conditions as parameterised by the outer scenario

This gives n*N valuation scenarios that are run through the underlying stochastic cash-flow model to generate the n tuples for the liability value (e.g., the BEL) to be approximated via LSMC.

Next, the LSMC calibration is performed to determine the best functional relationship between liability value and risk drivers as indicated by the data. The resulting function is used to evaluate the liabilities on a large set of real-world realisations of the risk drivers to obtain the distribution forecast and related figures such as SCR, quantiles, or expected shortfall.

By their nature, certain proxy models tend to have *black box* characteristics. That is to say, the user inputs a large set of data and the model outputs a functional relationship that incorporates complex effects of and interactions between risks or even just a distribution forecast. Therefore it is important to emphasise that each step of the LSMC process can and must be validated to give comfort and trust in the results. In the case study section below, we discuss different ways of validating LSMC results.

FIGURE 7: OVERALL LSMC PROCESS

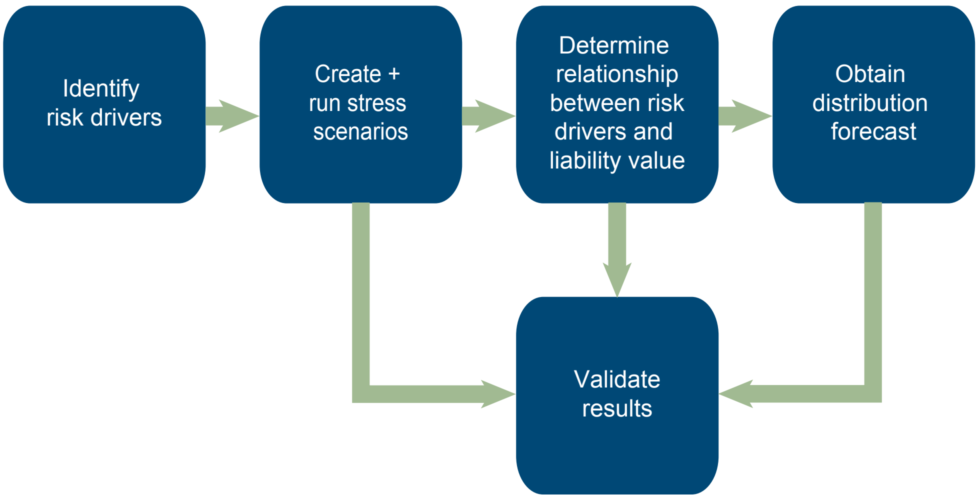

5.3.1 HOW TO CHOOSE RISK DRIVERS

Choosing a proper set of risk drivers and their respective parameterisations is an essential task for any kind of proxy modelling technique. The most vital point of this step is to find the right balance between including all of the risks that are relevant, while not overloading the exercise with thousands of potential risks that might be considered. Therefore, dimension reduction techniques such as, for example, a principal component analysis (PCA) of interest rates can help keep the number of risk drivers low.

LSMC is by no means restricted to market risk drivers but can also include insurance risks such as lapse, mortality, longevity, expenses or even management actions parameters (see the next section for details), because these can be included in a polynomial just the same as market risk drivers. In summary, LSMC can cover all those risk drivers which are accounted for by the underlying stochastic cash-flow model.

The number of relevant risk drivers varies from one application to another. While three risk drivers could well be enough for a simple valuation, upwards of 30 risk drivers could be relevant for a large-scale multi-currency situation.

5.3.2 DYNAMIC MANAGEMENT ACTIONS AND DYNAMIC POLICYHOLDER BEHAVIOUR

The LSMC approach aims at deriving a polynomial relationship between risk drivers and their effects on the liability value. Hence, by its nature LSMC automatically covers effects caused by dynamic management actions (DMA) or dynamic policyholder behaviour (DPB) because these effects are contained in the liability data used for the LSMC fitting.

More generally, one could also ask whether or how the LSMC process could allow for different sets of assumptions concerning DMA and DPB. On the one hand, if the LSMC calibration is carried out under a certain parameter set of DMA and DPB, then the LSMC results cannot be expected to remain appropriate under a different parameter set. On the other hand, it is possible to incorporate any key parameter of DMA or DPB into the risk driver list, if the LSMC polynomial has to act as a liability proxy under a range of different values of this parameter.

On a note of caution, while this approach can allow for possible changes of an important parameter, it clearly is not feasible to allow for changes in a large range of model parameters or to deal with discontinuities, e.g., dynamic purchases of a certain derivative instrument.

5.3.3 GENERATION OF OUTER SCENARIOS

The generation of the outer scenarios involves the following steps:

- For each risk driver X_i, we identify an interval $[a_i, b_i]$, $i=1,\ldots, m$ of values under consideration. Therefore we have to make sure that these individual intervals are large enough to cover all possible events we might experience with a certain probability, but simultaneously are not too wide; we want to restrict our focus to a relevant range of possible values.

  – The risk driver space under consideration is the multidimensional interval $[a_1, b_1]x\ldots x[a_m, b_m]$.

- The n outer positions are spread evenly over $[a_1, b_1]x\ldots x[a_m, b_m]$.

  – We do not include any real-world view in the process of choosing outer points. This would introduce a certain bias in the sense that the LSMC estimates will be particularly good and reliable for those risk driver realisations that appear with high probability, but display a significant loss of quality for rather unlikely risk driver realisations.[1]

  – The resulting liability function is based on information between risk drivers and liability values that have been evenly spread over all possible positions within the range under consideration and hence display a homogeneous level of quality even for tail events of the risk drivers.

1 Note that the tail realisations are the ones of most interest when it comes to SCR estimation.

FIGURE 8: DISTRIBUTION OF OUTER SCENARIOS

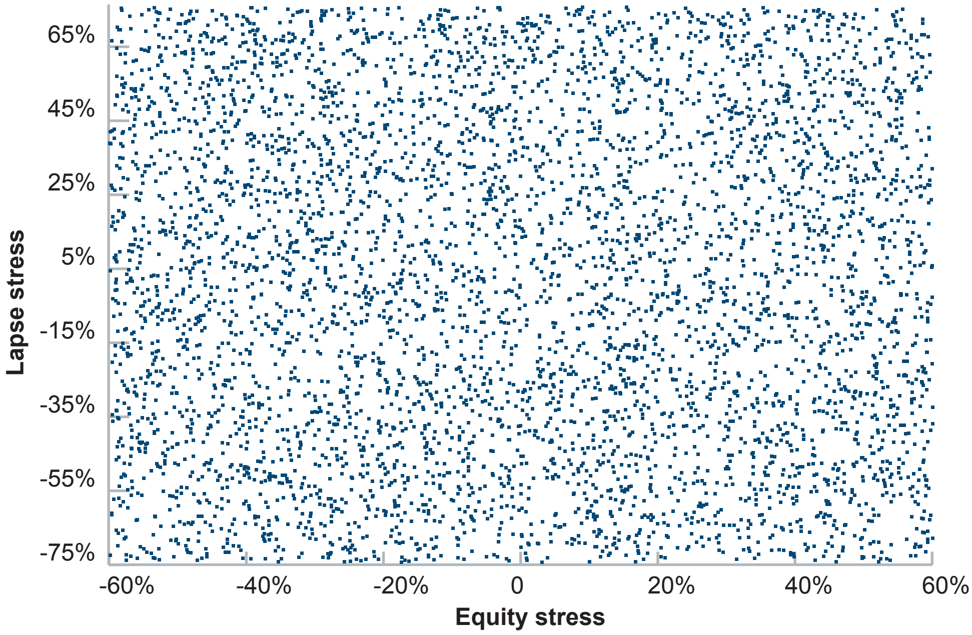

5.3.4 GENERATION OF INNER SCENARIOS: BRUTE FORCE OR BOOTSTRAPPING

Each of the outer scenarios considered in the previous section reflects different market conditions (such as a new initial yield curve or different volatility assumptions), which typically require a re-calibration of the economic scenario generator (ESG) underlying the generation of the inner scenarios at each point in time. This can pose significant challenges. However, there are a number of possible solutions to address this issue, including the following:

- Set up an automated solution for the ESG: In order for this to work properly, the ESG first needs to be provided with all of the outer scenarios that are of interest. For each of the scenarios, the ESG then needs to re-calibrate to the new conditions at each time step for which inner scenarios are required and generate the desired number of inner scenarios. This process is then repeated for each individual outer scenario. In practice, this may be achieved by using integrated scripting features of the ESG or an application programming interface that accesses the ESG.

- Apply the reweighted bootstrapping technique described in Hörig et al. (see References section below), which only requires a set of base valuation scenarios and transforms it into a large set of LSMC calibration scenarios. This is the approach that has been adopted for the purposes of this case study.

5.4 COMMENTS ON USE OF LSMC

5.4.1 ADVANTAGES OF LSMC

The LSMC proxy modelling approach displays the following strengths:

- LSMC involves only a limited amount of expert judgment as the statistical model (i.e., the polynomial and its coefficients that are used as proxy function) may be selected in an automated and objective way by using well-known statistical criteria.
- LSMC deals extremely efficiently with the data in the sense that it makes the most of the available data.
- Having the opportunity to derive confidence intervals for the LSMC liability value estimates gives valuable information on the reliability of the resulting estimates.
- Overall, the approach is generic enough to have real-world as well as risk-neutral application.
- Convergence of LSMC estimates against actual values can be proved mathematically.
- The approach can also include non-market risk drivers, such as changes to mortality rates, lapse assumptions, or management action parameters, because polynomial terms can be added for each of these risk drivers.
- Certain proxy modelling approaches can lead to instability, in the sense that small changes in the input data might lead to large changes in the resulting proxy functions and their parameters. This problem typically arises when companies calibrate their proxy models on an annual or quarterly basis. While one would expect that subsequent proxy models look quite similar, it is often the case that the resulting models are completely different (in terms of functions and coefficients involved). This property is clearly unsatisfactory because a gradually changing liability portfolio should not lead to proxy models that are completely discontinuous over time. Stable models with reasonable parameters allow companies to build some intuition about how the business should behave as a function of the risk drivers.
  - Our implementation of the LSMC approach (for the purposes of this case study) yields robust proxy functions by using a combination of special basis functions and outer calibration scenarios that are uniformly distributed over the space of possible risk driver values (so-called orthogonal outer scenarios).[2]

5.4.2. IMPLICIT ASSUMPTIONS AND LIMITATIONS

One of the most significant assumptions underlying the LSMC technique is that the error variable (ε) is normally distributed. Illustratively this means that when simulating the liability value with only a limited number of valuation scenarios, the difference between this rough estimate and the actual value is symmetrically distributed and tends to be negative or positive with the same probabilities. While this property holds asymptotically (when using more and more valuation scenarios) by virtue of the central limit theorem, it might be violated for some valuation scenarios once the liability value itself possesses some limitations in the sense that it is bounded from above or below.

Furthermore, the main paradigm of using polynomial proxy models is that the liability value to proxy can be properly approximated by polynomial terms. Once the underlying model displays severe discontinuities–for example, based on certain management actions–this may result in a certain degree of loss of accuracy. However, while such discontinuities can appear when dealing with a single policy or scenario, they are often smoothed out within a larger portfolio. Finally, the impact of management rule discontinuities is typically smoothed over time within a net present value such as a Solvency II BEL.

2 Note that such techniques can also be applied to other proxy approaches such as replicating portfolios in order to get robust estimates.

6 CASE STUDY

6.1 INTRODUCTION TO VARIABLE ANNUITIES

In this paper we focus our attention on the analysis of capital requirements for typical variable annuity (VA) products currently sold in the market.

The term VA can be applied to a wide range of product types. Generally it is used to describe products which are unit-linked in nature but with investment and/or insurance guarantees provided to the policyholder. A detailed analysis of the type of products in question is contained in 'Variable Annuities,' an Institute of Actuaries research paper (see References section below). Here is a high-level summary of the various types of guarantees that are commonly available.

The guarantees provided on VAs generally fall into four categories as follows:

- Guaranteed minimum death benefit (GMDB). In its most basic form, this guarantee represents a return of the premium invested on death if the underlying unit-linked account value is less than the premium invested.

- Guaranteed minimum accumulation benefit (GMAB). This guarantee can take similar forms to the GMDB except that it applies on survival for a specified period of time. There may be multiple times during the duration of the policy when the guarantee applies.

- Guaranteed minimum income benefit (GMIB). This guarantee provides a regular income stream from a specified future point in time, typically in the form of an annuity.

- Guaranteed minimum withdrawal benefit (GMWB). This guarantee offers a level of withdrawals from the underlying unit-linked fund value even if the account value runs out. The guaranteed income stream can apply for either a fixed period of time or for life. The 'for-life' version is commonly referred to as a GMWBL.

Note that GMIBs and GMWBs are normally associated with a pension arrangement so there may be a pre-retirement accumulation phase before the guarantees kick-in.

The guaranteed amount is generally known as the benefit base (BB). The basic form of VA guarantee involves the return of the BB (or payment of an income stream based on it) to the policyholder, contingent on some event (death, survival, or maturity). The product design can involve certain variations though; for example, the BB can accumulate, or *roll up*, at a defined rate each year. It can also be locked in at higher amounts, depending on the performance of the underlying unit-linked account value (something commonly known as a *ratchet* feature).

6.2 SAMPLE PRODUCT DESIGN

For the purposes of this case study we considered two individual model point policies, one representing a GMAB policy and the other a GMWB policy. The policies can be considered to have just come into force and the guarantees are at-the-money. The details of the model points are set out in Appendix A.1.

Our case study analysis of these model points captures a wide range of features and risks of VA products including dynamic policyholder behaviour, a range of guarantee types, fund performance risk stemming from equity and bond returns, and interest rate risk.

6.3 LAPSE ASSUMPTIONS

The chance of a policyholder surrendering a VA policy depends on how the account value performs relative to the guarantee level. Policyholders are less likely to surrender a policy if the guaranteed benefits are very valuable.

Therefore, we have used a dynamic lapse model to capture some sensitivity of lapses to fund performance. This sample dynamic formula is included in Appendix A.1. The structure and calibration of the dynamic formula is purely for illustrative purposes in order to demonstrate the LSMC approach.

6.4 MODELLING SIMPLIFICATIONS

We have only considered two independent model points representing a GMAB case and a GMWB case. We also did not model a portfolio of policies, but nevertheless this can be dealt with adequately using the LSMC approach.

The pricing of the guarantee charges was set to give a reasonable level of profitability to an insurer under conditions pertaining at the end of 2012 but we have not carried out a detailed pricing exercise. The use of the two specific model points is for illustrative purposes only in order to help to demonstrate the LSMC approach.

Both model points are new business model points, so when we consider cash flows at time t=1, for example, this represents the first policy anniversary.

6.5 SETUP

The case study was performed using GMAB and GMWB product versions, as outlined above. Given the nature of these two products, we've considered the following risk drivers:

- Interest rates: Interest rate risk refers to the risk arising from a change of the interest yield curve. Hence, it is not a one-dimensional risk but depends on the structure of the whole yield curve. Historical time series show that interest rate yields for different maturities are not independent but move with a high level of mutual dependency. Hence, considering each maturity as an independent risk driver would not be reasonable because most joint combinations of these risk drivers would never be experienced in reality.

  Therefore, we apply a principal component analysis (PCA), which is a statistical method that (in this context) aims at reducing the number of dimensions of the multivariate time series. The starting point for the PCA is a historical time series of EUR swap rates ranging from 31 March 2003 to 31 March 2013. The implementation of the PCA method follows that set out by I.T. Jolliffe (see References section).

  Based on the PCA, we see that the first principal component (PC1, the level of the curve) explains 94% of the yield curve variation during the past 10 years and that the second principal component (PC2, the slope of the curve) explains another 4%. Hence, using the first two PCs as risk drivers explains 98% of the historical yield curve variation.[3] Therefore, it is reasonable to focus on these two PCs as risk drivers for interest rates.

[3] A PCA performs an Eigenvector/Eigenvalue decomposition of the covariance matrix of historical yield curves. Each Eigenvector refers to one PC. The importance of the Eigenvector is determined by the magnitude of the corresponding Eigenvalue. The sum of the first two Eigenvalues (the weights of the first two PCs) equals 98% of the sum over all Eigenvalues.

For the generation of different initial interest yield curves, the general approach that we took was to simultaneously stress the first two principal components of the yield curve within reasonable ranges.[4] Given the current very low interest rate environment we used a non-market yield curve as part of the calibration process (see Appendix A.2 for a comparison of this yield curve with a market yield curve). The aim of this procedure is to perform sensible yield curve stresses by stressing the most important building blocks of the full yield curves, i.e., the first two principal components, which implicitly stresses all maturities. By doing so we avoid the challenge of determining 'realistic' stresses for each maturity itself, because it is very unlikely that, for example, the 7-year rate goes up by 2% while simultaneously the 8-year rate drops by 2%.

- Equity return: Initial equity return was stressed between -60% and 60% to cover a reasonable range of equity movements.

- Equity volatility: Initial equity volatility, i.e., the implied volatility of an at-the-money (ATM) call option on equity with maturity in 10 years was 25%. We've stressed this volatility to levels between 10% and 60%.

- Bond fund return: Initial bond fund return was stressed between -30% and 30%.

- Lapse rates: Lapse assumptions have been stressed with multiplicative factors between 0.25 and 1.75.

For each product, cash flows were simulated across a total of 50,000 scenarios, comprising 5,000 outer scenarios with 10 inner valuation scenarios per outer scenario. There is no specific required number of scenarios to use to reach a certain level of precision (noting that we know from Section 5.2.4 that the proxy model converges towards the true values once we continue to increase the number of inner and outer scenarios and allow for polynomial terms of higher order). However, as part of the assessment of the quality of our proxy model we will see that 50,000 scenarios leads to a good fit.

For the purposes of the case study, the variable of interest was the BEL at time t=1. Running the 50,000 calibration scenarios through the liability valuation model resulted in 5,000 different values for the BEL, each one derived as the mean over 10 inner valuation scenarios.

In order to test the capability of the LSMC proxy methodology in dealing with dynamic actions (either management actions or policyholder behaviour), we carried out the fitting procedure twice: first including dynamic lapse behaviour and then without the dynamic element of lapse behaviour. The figures and graphs in the following sections relate to the case, which includes dynamic policyholder behaviour (which we consider to be the more relevant in the case of VA business because dynamic policyholder behaviour is a key element of actuarial projection models for this type of business). However, when performing a quantitative assessment of the quality of the LSMC proxy models in our case study (see Sections 7.1 and 7.2) we also include results for the case without dynamic policyholder behaviour in order to compare the explanatory quality of the proxy model in both cases.

The LSMC proxy model was then calibrated in accordance with the steps described in Sections 5.1 to 5.3.

4 When determining a reasonable stress range, we restricted our focus in order to avoid the possibility of negative rates occurring at any maturity.

7 DISCUSSION AND VALIDATION OF LSMC CALIBRATION

Having performed the LSMC calibrations for the GMAB and GMWB products under consideration, we now consider the quantitative and qualitative assessment of the LSMC proxy models. In doing so we consider three layers of validation:

- Qualitative assessment of the two-dimensional and three-dimensional function plots for different risks, that is:

  — Do the function plots reflect the observations that were used as inputs for the calibration?
  — Do we see monotonicity in the plot where this is expected?
  — Are the level, slope and curvature of the graphs as expected?
  — Do three-dimensional plots display interactions between the main risk drivers?

- Consideration of confidence intervals for LSMC estimates

- Comparison of LSMC estimates with actual BEL values derived via a full Monte Carlo simulation with 2,000 valuation scenarios for each risk driver position

The following sections first cover the results of the GMAB product then the results of the GMWB product.

7.1 GMAB LIABILITIES: QUALITATIVE ASSESSMENT AND CONFIDENCE INTERVALS

Figure 9 illustrates the starting point for the LSMC fitting as seen from the bond fund return risk driver. Each blue dot refers to the estimate for the BEL for one outer scenario, i.e., the mean value over 10 inner valuation scenarios per outer scenario. We can clearly see that these 5,000 estimates display quite a high level of variance[5] but still show the clear tendency to decrease once the initial bond fund return is increasing, which is to be expected.

FIGURE 9: GMAB SCATTERPLOT

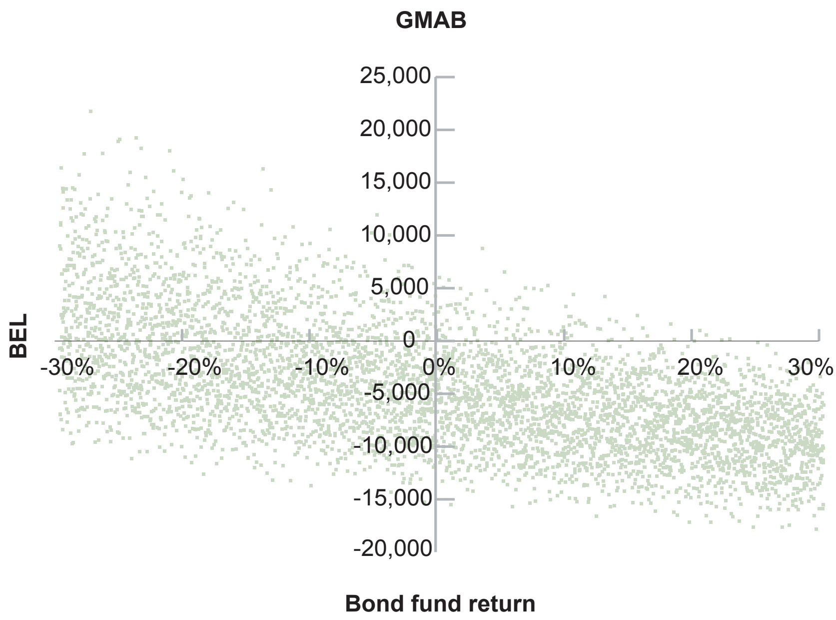

[5] Note that the extent of the variation of the observations is also due to the variation in the other risk drivers for each observation.

Figure 10 shows a two-dimensional plot of the LSMC proxy function for the bond fund return, which shows that the curve is smooth through the noisy input data from Figure 9.

FIGURE 10: GMAB LIABILITY VALUE: 2D PLOT FOR BOND FUND RETURN

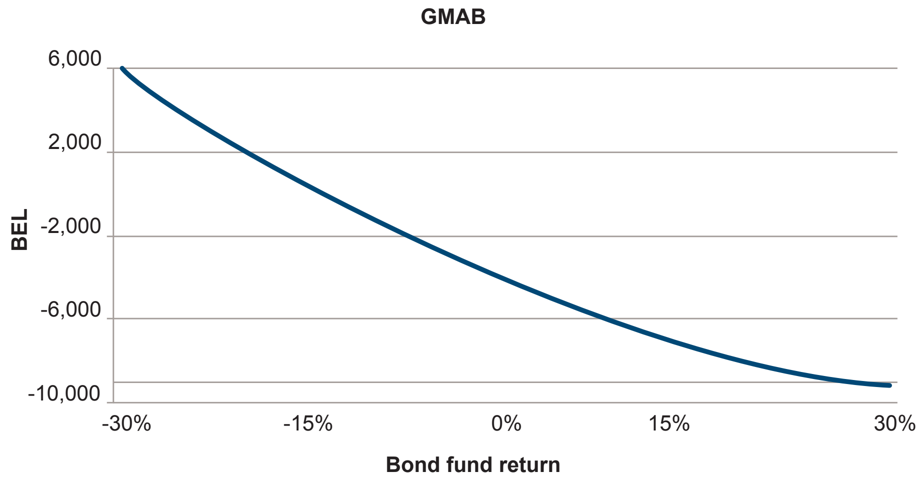

The following two-dimensional plots (Figures 11 to 14) show the shape of the LSMC function for each of the individual risk drivers while all other risk drivers are kept at their unstressed level from the calibration process (i.e., the plots show the marginal sensitivities for each individual risk driver):

- Figure 11: This plot of BEL versus PC1 shows decreasing BEL values once the PC1 increases, i.e., once interest rates go up. This is caused by a number of effects:[6]
  - Higher interest rates lead to higher discounting rates, which decrease BEL values.
  - Higher initial interest rates also increase the drift of the equity position, which again decreases BEL values.
  - Higher interest rates can also lead to lower lapses (as well as possibly lower utilisation of income benefits in the case of GMWB), depending on the form of the dynamic policyholder behaviour assumptions (i.e., if interest rates form part of the dynamic function).

    Note that the PC1 value which refers to the position of the base (market) yield curve is actually close to PC1=-0.5 as opposed to corresponding to PC=0. This follows from the calibration of the proxy model (as discussed in Section 6.5 above), during which a modified yield curve was used to determine the polynomials. This is because the application of symmetrical interest rate stresses to the current market yield curve would otherwise inevitably lead to negative initial yields given current conditions.

- Figure 12: This plot of BEL versus equity performance shows that BEL values decrease as equity returns increase. However, the relationship is non-linear. This is discussed further below.

- Figure 13: The liability value increases as equity volatility increases, which is primarily due to the fact that the cost of guarantee claims will increase with higher equity volatility.

- Figure 14: Higher lapse rates lead to decreasing liability values. Normally, one would expect that higher lapse rates would lead to lower liability values for an at-the-money guarantee (as, in this case, charges arise before the benefit payments and so higher lapses have a proportionately greater impact on the benefits than the charges). However, higher lapses also erode margins on the underlying product, leading to a less negative BEL (i.e., a higher liability value). In this case, the impact on the guarantee outweighs the impact on the base product margins. While this plot is reasonable in this case it is important to note that this effect may depend on many factors, such as changes in the economic environment and the passage of time.

It is worth highlighting that the two-dimensional plots of risk drivers with particularly strong influence on the BEL (such as equity performance, equity volatility, and bond fund return) typically show non-linear behaviour, while the plots of less-severe risk drivers (such as interest rates and lapse) display linear behaviour. This indicates that the more severe risk drivers trigger some asymmetrical events, which cause non-linear shapes, while the less severe ones have symmetrical influence on the liability values.

Consider equity performance as an example. Equity performance will affect the margins associated with the unit-linked account value. This element will be reasonably linear. However, because of the existence of the ratchet feature on the VA guarantee, increases in equity values will drive benefit ratchets, meaning that guarantees will remain at-the-money even if investment performance is good. On the other hand, falls in equity values will result in guarantees biting. This generates non-symmetrical behaviour.

[6] Note that there is also a counteracting effect here: higher interest rates will reduce bond values, which will increase the BEL.

FIGURE 11: GMAB LIABILITY VALUE: 2D PLOT FOR PC1

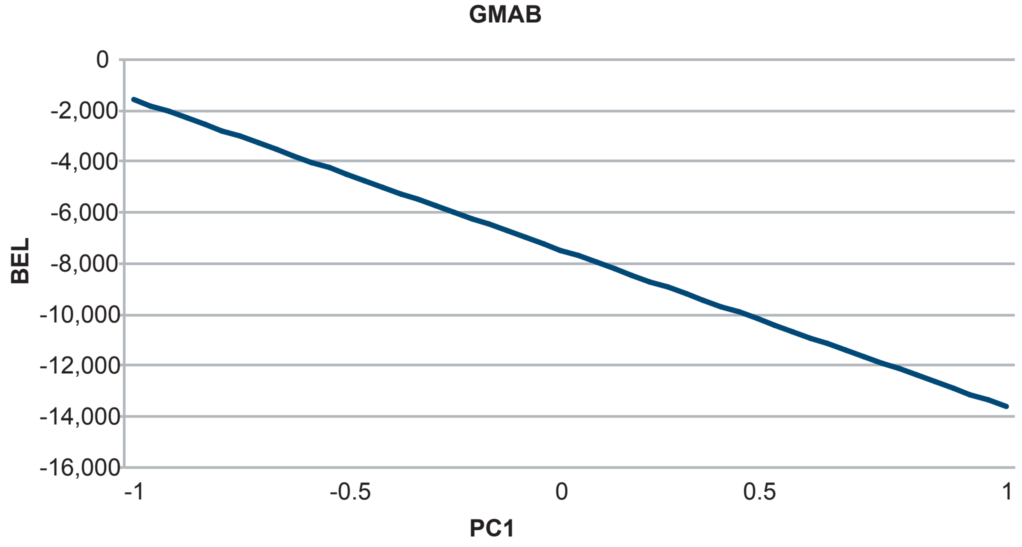

FIGURE 12: GMAB LIABILITY VALUE: 2D PLOT FOR EQUITY FUND RETURN

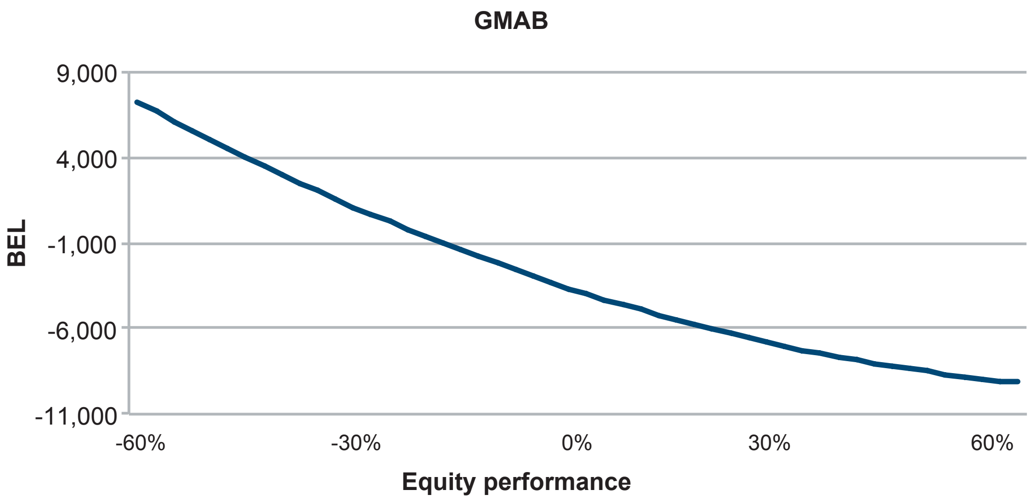

FIGURE 13: GMAB LIABILITY VALUE: 2D PLOT FOR EQUITY VOLATILITY

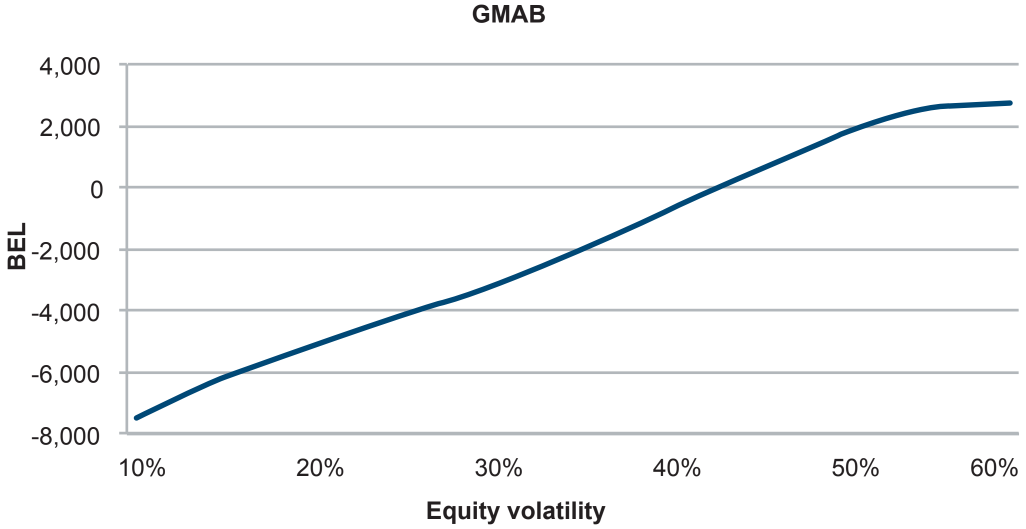

FIGURE 14: GMAB LIABILITY VALUE: 2D PLOT FOR LAPSE

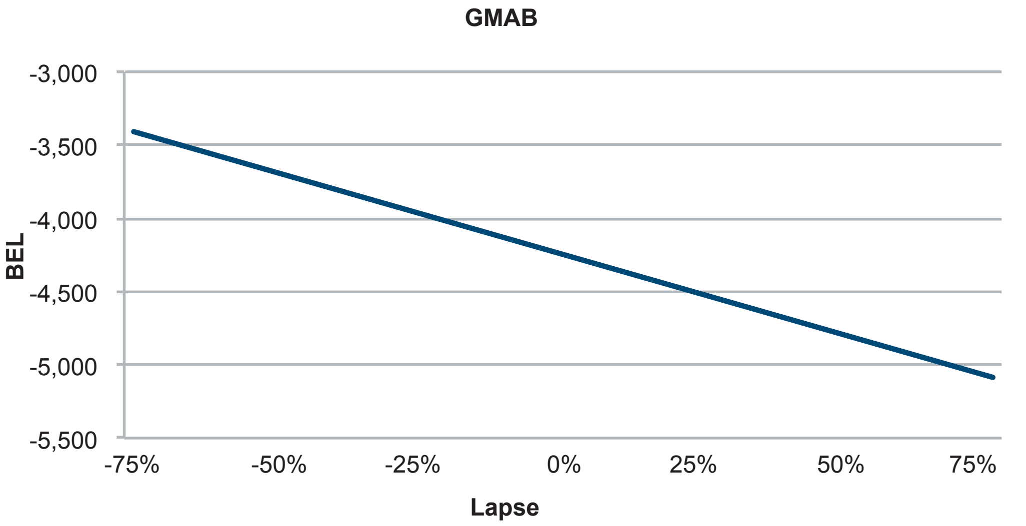

The three-dimensional plot in Figure 15 illustrates the joint behaviour of the BEL in the two risk drivers PC1 and lapse. Similar plots can easily be constructed for any pair of risk drivers relative to the BEL. The structure of the surface in Figure 15 clearly indicates that there are strong interactions between interest rates and lapse rates.

At high interest rate levels, the higher the lapse rate the higher the BEL. This is because higher interest rates cause the guarantee to move out-of-the-money. Hence, the guarantee itself becomes an asset (i.e., the present value of charges exceeds the present value of claims). As lapse rates increase, the value of this asset reduces (i.e., the BEL increases). At low interest rate levels, the opposite occurs.

While we know from the projection models for the GMAB product that there are interactions between interest rates and lapse rates, Figure 15 underlines that they have been detected and captured in a sensible way by the LSMC proxy model.

FIGURE 15: GMAB LIABILITY VALUE: 3D PLOT FOR PC1 AND LAPSE

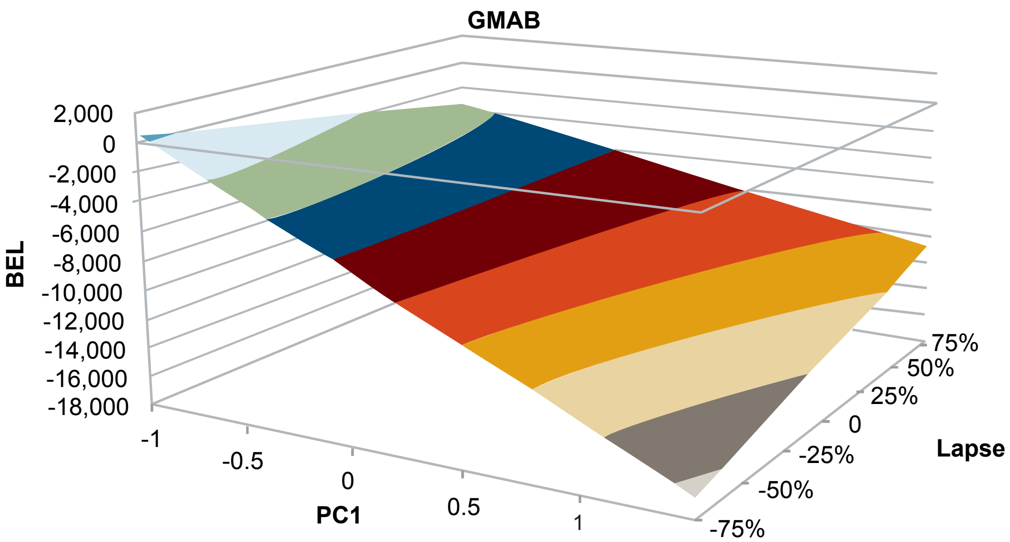

Figure 16 displays a two-dimensional plot for the liability value with varying initial levels of equity performance (blue line) as well as upper and lower 90% confidence intervals (green lines). The confidence intervals have been derived using the jackknifing procedure (see Appendix A.3 for further details). This involves fitting proxy models to subsets of the full set of fitting data. Any such new proxy model serves as a random realisation of the main proxy model and we therefore get a variety of liability value estimates for any risk driver position by evaluating each of these proxy models. Based on these realisations we have derived an empirical distribution of any liability value estimate of the proxy model, which in turn allows for the assessment of the corresponding confidence intervals. Hence, these confidence intervals span the range which contains any particular LSMC estimate with probability 80% and indicate how much randomness is left in the corresponding estimates. The confidence intervals seen here are reasonable in the sense that they are quite narrow and display a very similar pattern to the LSMC curve itself. However, this is a subjective assessment and requires a certain amount of expert judgment.

FIGURE 16: GMAB LIABILITY VALUE: CONFIDENCE INTERVALS FOR EQUITY PERFORMANCE

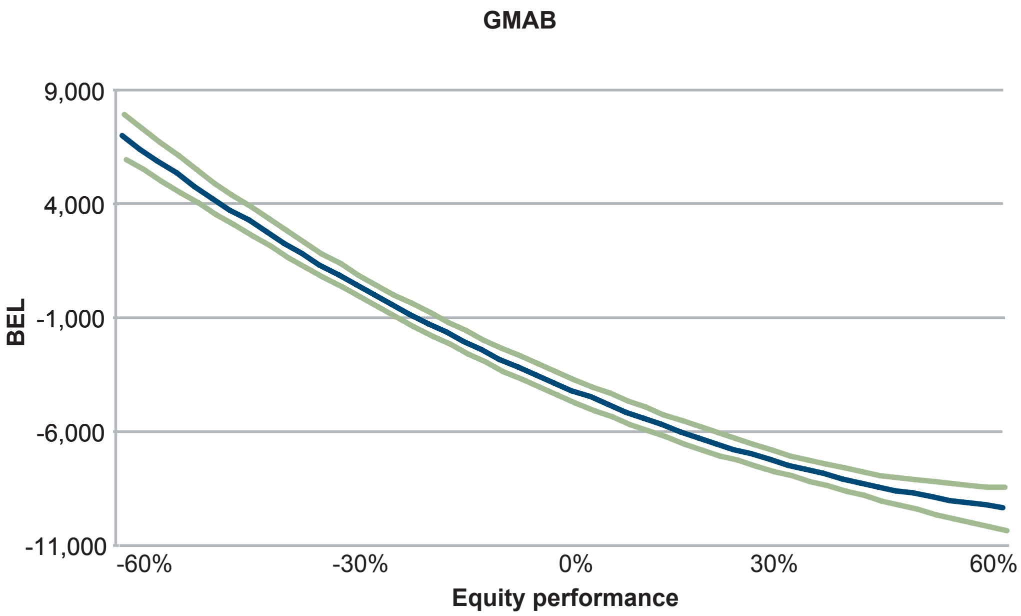

Turning now to a quantitative validation of the model, we compare the LSMC estimates with their corresponding actual values using the full liability cash-flow model. This involves evaluating the BEL for some risk driver combinations with a full set of 2,000 valuation scenarios each and comparing those values to the corresponding LSMC estimates. Figure 17 shows the plot of liability value against the bond fund return (while the other risk drivers remain unchanged) in blue as well as seven actual BEL values for different bond fund returns (red dots). Figure 17 also shows the corresponding upper and lower 90% confidence intervals (green lines). We see that the actual values are close to the proxy function values, display the same overall shape, and fall within the bounds of the LSMC confidence intervals. For a more detailed assessment of these validations, refer to the table in Figure 25 in Section 7.3.

FIGURE 17: GMAB LIABILITY VALUE: CONFIDENCE INTERVALS AND ACTUAL VALUES FOR BOND FUND RETURN

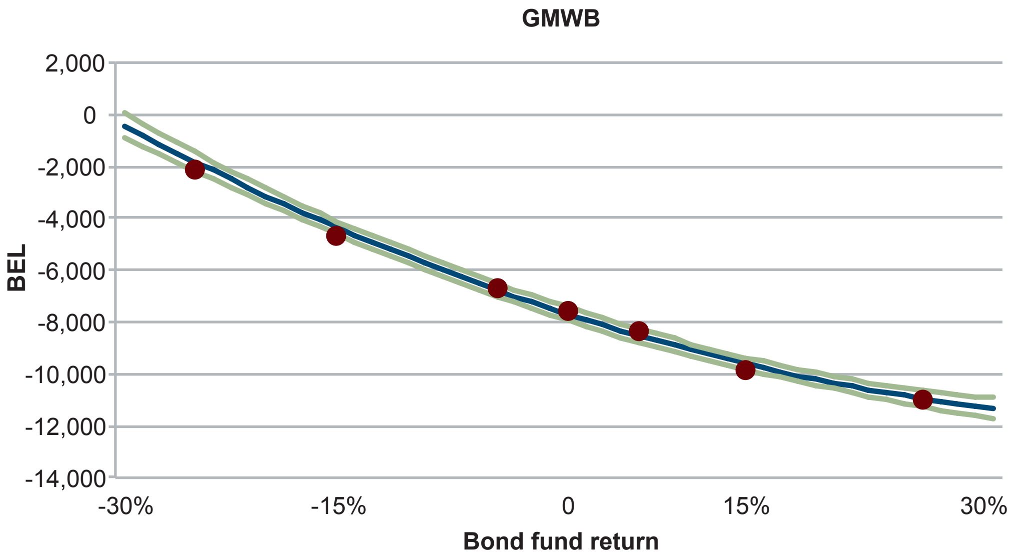

7.2 GMWB LIABILITIES: QUALITATIVE ASSESSMENT AND CONFIDENCE INTERVALS

This section considers the GMWB product. As with the GMAB product, we start with a scatterplot of each risk driver versus the BEL. Figure 18 shows 5,000 BEL estimates for different levels of equity volatility. Each blue dot refers to the estimate of the BEL for one outer scenario, i.e., the mean value over 10 valuation scenarios per outer scenario. Just as in Section 7.1, there is a high degree of noise in these initial estimates.

FIGURE 18: GMWB SCATTERPLOT

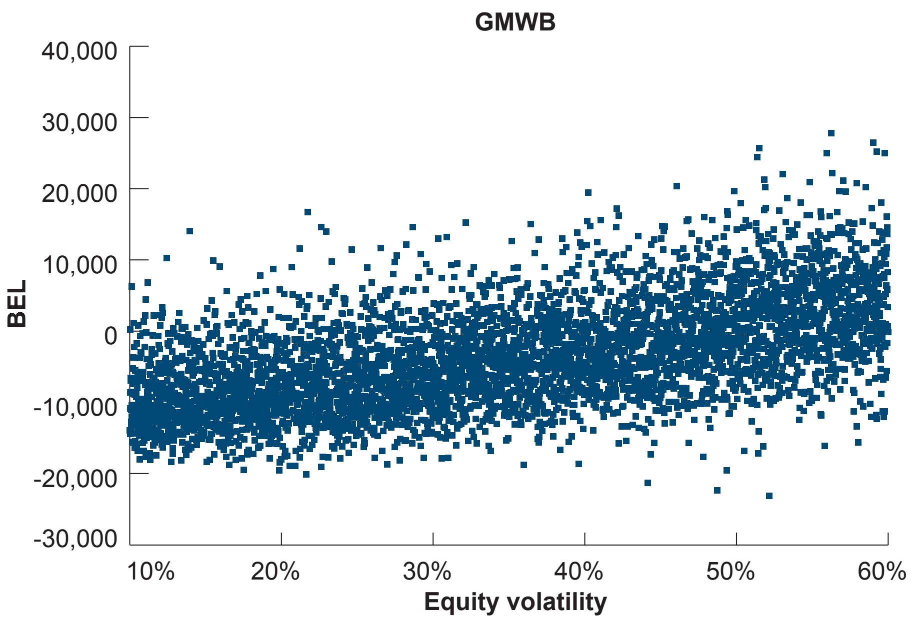

Figure 19 shows a two-dimensional plot of the LSMC proxy function for equity volatility. We can see that the curve is smooth through the noisy input data from Figure 18 and maps the best estimate of the true relationship between equity volatility and liability value. As in Section 7.1, the value of the BEL increases as the equity volatility increases, which is primarily due to the fact that the cost of guarantee claims will increase with higher equity volatility (as the guarantee is more likely to bite and guarantee ratchets are more likely to occur).

FIGURE 19: GMWB LIABILITY VALUE: 2D PLOT FOR EQUITY VOLATILITY

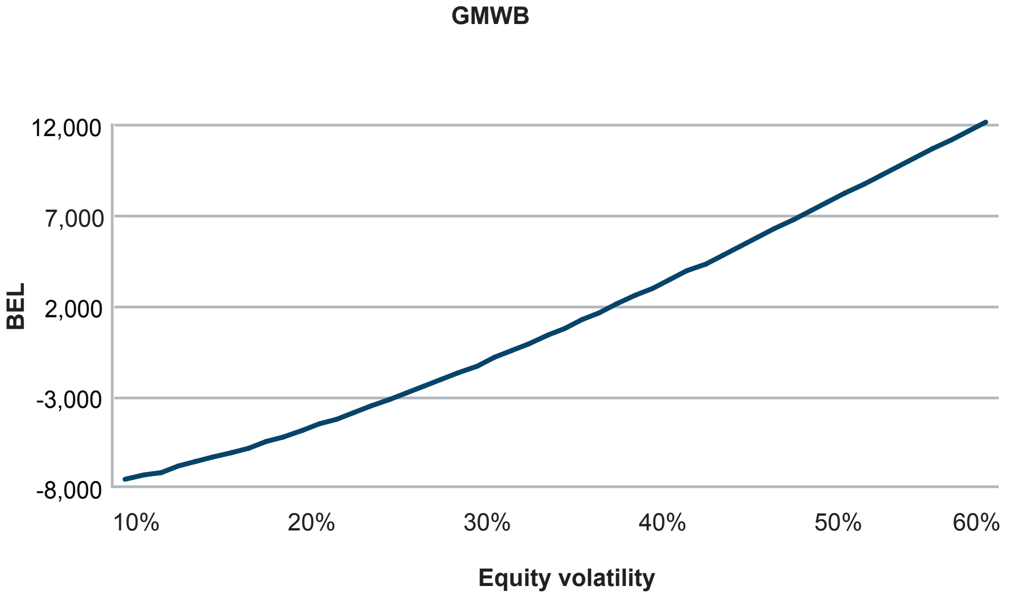

The following sample two-dimensional plots (Figures 20 and 21) show the shape of the LSMC function for PC1 and equity return risk drivers, respectively, while all other risk drivers in each case are kept at their unstressed levels from the calibration process (i.e., the plots show the marginal sensitivities for each individual risk driver):

- Figure 20: The BEL versus PC1 plot shows decreasing BEL values once the PC1 increases, i.e., once interest rates go up. The same drivers apply here as for the GMAB product (see Figure 11 in Section 7.1 above), although the relationship is much less linear. This is due to the greater impact of interest rates on the value of the GMWB BEL relative to the GMAB. In Figure 20, as interest rates rise (PC1=0 to PC1=0.5) the value of the BEL falls. This is caused by the reduction in the present value of charges being outweighed by the reduction in the present value of benefits. The duration of the benefits is much longer than the duration of the charges. Hence, as rates continue to rise (PC1=0.5 to PC1=1.0) a convexity effect starts to kick in, which prevents the value of the BEL decreasing further.

- Figure 21: BEL values decrease once the initial equity position increases, which is a reasonable reaction. In such a situation, the GMWB claims will reduce and the fund charges will increase so it makes sense that the BEL becomes more negative.

FIGURE 20: GMWB LIABILITY VALUE: 2D PLOT FOR PC1

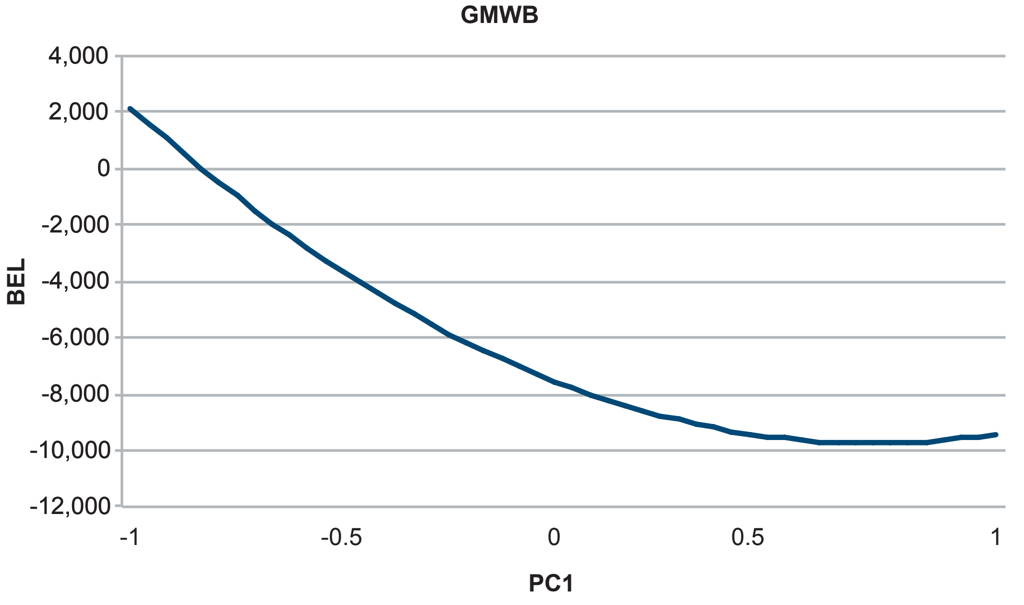

FIGURE 21: GMWB LIABILITY VALUE: 2D PLOT FOR EQUITY PERFORMANCE

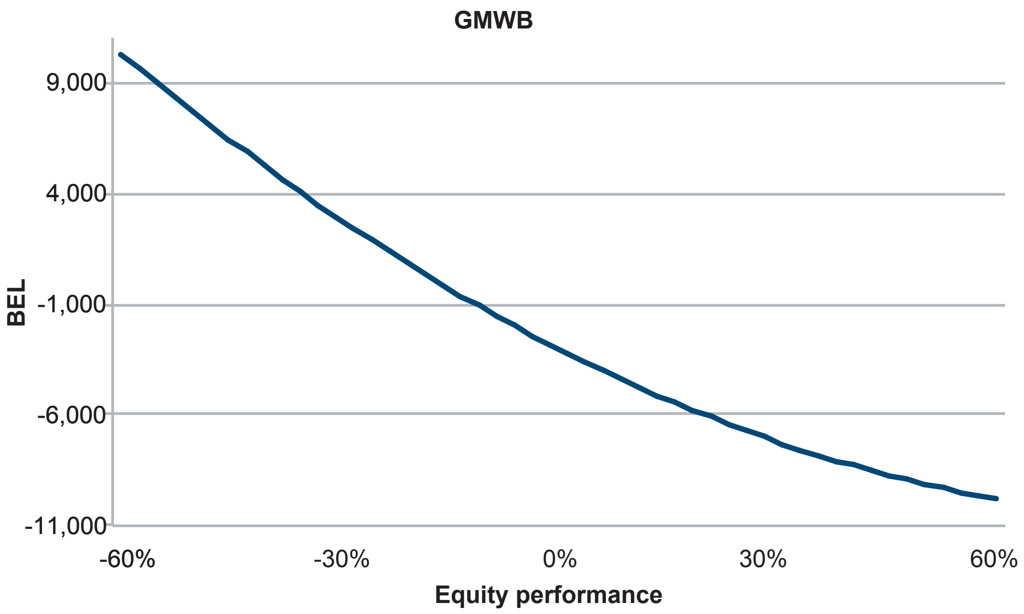

Similar plots can be constructed for the remaining risk drivers (as shown in Section 7.1 for the GMAB product).

Next we turn to an assessment of potential interactions and whether they have been captured by the LSMC proxy model. The three-dimensional plot in Figure 22 illustrates the joint behaviour of the BEL in accordance with the two risk drivers PC1 and lapse. Just as for the GMAB product in Section 7.1, we see that the evolving surface clearly indicates the presence of interactions between interest rates and lapse rates. These interactions are broadly similar to what is observed in the case of the GMAB above in that, at high interest rate levels, the higher the lapse rate the higher the BEL, while at low interest rate levels, the opposite occurs.

FIGURE 22: GMWB LIABILITY VALUE: 3D PLOT FOR PC1 AND LAPSE

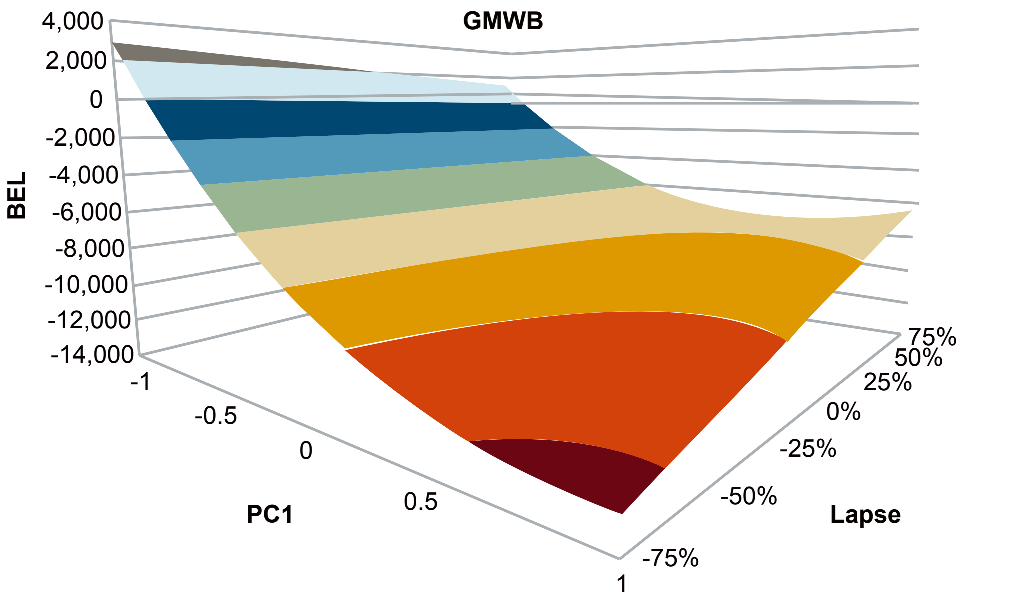

Figure 23 shows a two-dimensional plot for the liability values (blue line) with varying equity volatility as well as upper and lower 90% confidence intervals (green lines). As described in Section 7.1 above (see Figure 16), the confidence intervals seen here are reasonable in the sense that they are quite narrow and display a pattern very similar to the liability plot itself.

FIGURE 23: GMWB LIABILITY VALUE: CONFIDENCE INTERVALS FOR EQUITY VOLATILITY

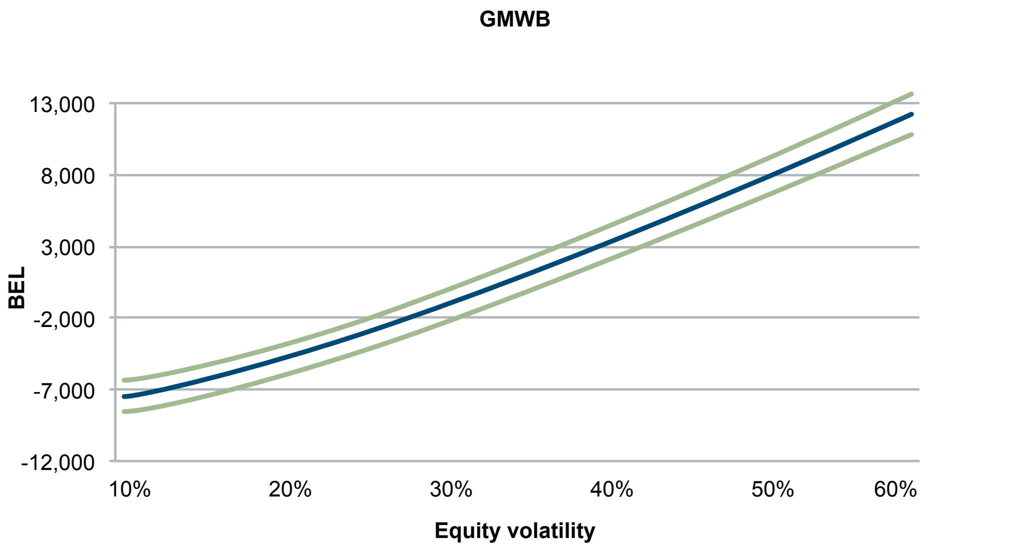

In the remainder of this section we perform a quantitative validation of our overall results. Figure 24 shows the plot of liability value (blue line) against the bond fund return (while the other risk drivers are fixed at a certain level). Seven actual BEL values for different levels of the bond fund return (red dots) are shown, as well as the corresponding upper and lower 90% confidence intervals (green lines). Here, we see that the proxy values are good estimators for the actual values and both display the same overall shape.

FIGURE 24: GMWB LIABILITY VALUE: CONFIDENCE INTERVALS AND ACTUAL VALUES FOR BOND FUND RETURN

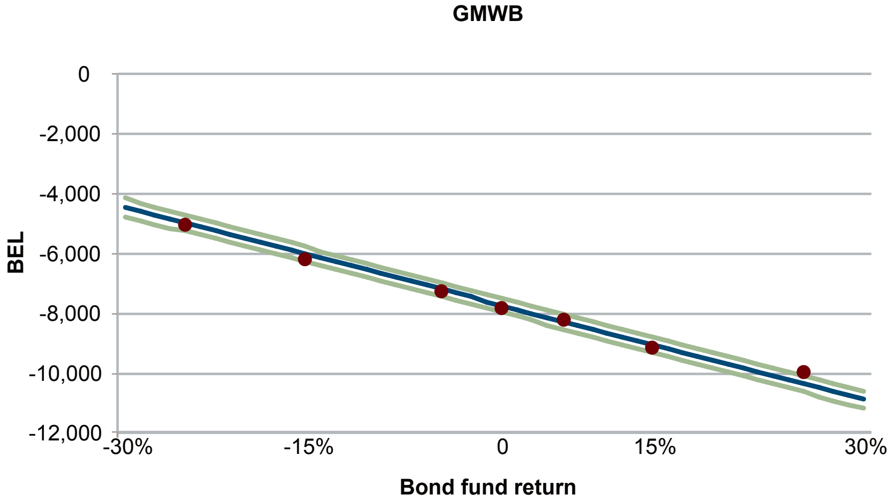

7.3 COMPARISON OF LSMC ESTIMATES WITH ACTUAL BEL VALUES

In this section we compare estimates prepared using the LSMC proxy model for a set of sample scenarios with their corresponding actual values using the full liability cash-flow model.

The table in Figure 25 compares the BEL estimates derived by the proxy model with actual values derived by Monte Carlo simulations for 37 different risk driver positions (the base one and six uniformly distributed ones per risk driver). In practice, it is preferable to also test combinations of stressed risk drivers, though we have not done so here (in the interests of simplicity).

FIGURE 25: VALIDATION DATA

		GMAB			GMWB		
Risk Driver	Stress Level	LSMC Value	Actual Value	Deviation	LSMC Value	Actual Value	Deviation
None	Base level	-7,582	-7,579	0.00%	-7,671	-7,806	1.70%
Interest Rates PC1	-0.9	-9,393	-9,105	3.80%	-9,046	-8,972	0.90%
	-0.6	-8,790	-8,680	1.40%	-8,669	-8,687	0.20%
	-0.3	-8,186	-8,175	0.10%	-8,210	-8,304	1.20%
	0.3	-6,979	-6,889	1.20%	-7,051	-7,178	1.60%
	0.6	-6,375	-6,094	3.70%	-6,350	-6,396	0.60%
	0.9	-5,771	-5,176	7.90%	-5,567	-5,439	1.60%
Interest Rates PC2	-0.9	-7,342	-7,545	2.70%	-7,462	-7,572	1.40%
	-0.6	-7,327	-7,556	3.00%	-7,532	-7,658	1.60%
	-0.3	-7,496	-7,567	0.90%	-7,602	-7,736	1.70%
	0.3	-7,496	-7,590	1.20%	-7,741	-7,870	1.70%
	0.6	-7,327	-7,601	3.60%	-7,811	-7,927	1.50%
	0.9	-7,342	-7,613	3.60%	-7,880	-7,977	1.20%
Equity Return	-50%	-1,224	-1,536	4.10%	-1,204	-1,421	2.80%
	-30%	-4,099	-4,357	3.40%	-4,114	-4,387	3.50%
	-10%	-6,532	-6,618	1.10%	-6,593	-6,790	2.50%
	10%	-8,522	-8,449	1.00%	-8,642	-8,720	1.00%
	30%	-10,071	-9,948	1.60%	-10,260	-10,310	0.60%
	50%	-11,177	-11,214	0.50%	-11,448	-11,667	2.80%

FIGURE 25: VALIDATION DATA – CONTINUED

		GMAB			GMWB		
Risk Driver	Stress Level	LSMC Value	Actual Value	Deviation	LSMC Value	Actual Value	Deviation
Equity Volatility	20.00%	-8,377	-8,291	1.10%	-8,735	-8,657	1.00%
	21.25%	-8,178	-8,056	1.60%	-8,488	-8,385	1.30%
	22.50%	-7,981	-7,818	2.10%	-8,228	-8,101	1.60%
	23.75%	-7,783	-7,338	5.90%	-7,956	-7,503	5.80%
	26.25%	-7,377	-7,096	3.70%	-7,375	-7,190	2.40%
	28.75%	-6,946	-6,853	1.20%	-6,749	-6,869	1.50%
Lapse*	0.60	-7,848	-7,879	0.40%	-8,285	-8,602	4.10%
	0.75	-7,748	-7,763	0.20%	-8,055	-8,278	2.90%
	0.90	-7,649	-7,651	0.00%	-7,825	-7,986	2.10%
	1.10	-7,516	-7,508	0.10%	-7,518	-7,638	1.50%
	1.25	-7,416	-7,406	0.10%	-7,288	-7,403	1.50%
	1.40	-7,316	-7,306	0.10%	-7,057	-7,188	1.70%
Bond Fund Return	-25%	-1,852	-2,079	3.00%	-4,991	-5,039	0.60%
	-15%	-4,428	-4,625	2.60%	-6,063	-6,238	2.20%
	-5%	-6,626	-6,690	0.80%	-7,135	-7,313	2.30%
	5%	-8,444	-8,390	0.70%	-8,207	-8,274	0.90%
	15%	-9,885	-9,800	1.10%	-9,279	-9,144	1.70%
	25%	-10,947	-11,000	0.70%	-10,351	-9,937	5.30%
Average				**1.90%**			**1.90%**

*Note: PCI=0 for the base level. * Stress value refers to the factor that is applied to base lapse rates.*

We see that the average deviations between the LSMC estimates and the actual values are 1.9% for both products. This indicates that the LSMC proxy model generally provides a good approximation for the behaviour of the true liability values over a large variety of economic circumstances and that we have achieved a certain level of convergence judging by these out-of-sample validations. Here, the convergence becomes visible by that fact that the overall level of deviation is quite stable for different kinds of economic stresses applied to the liabilities. Therefore, we can expect a similar level of deviation in the SCR, for example, since the economic conditions underlying the SCR calculation will be within the bounds of the range of conditions for which we have performed the validation. Having assessed these deviations allows us to make a judgment as to whether or not the achieved level of convergence is sufficient. By increasing the overall number of simulations we provide a simple framework to further increase the level of convergence. Of course, this is a subjective assessment, and the level of deviation which may be considered acceptable will vary from one situation to the next.

Deciding upon whether or not the deviations between the liability value estimates derived from the proxy model and the actual liability values are at an acceptable level is a non-trivial task, for the following reasons:

- Choice of an error measure: For the purposes of this case study, we compare the absolute value of the deviations to the base liability value. This error measure can be misleading once the base value is close to zero because the resulting relative deviations are artificially magnified in this case. Furthermore, if the liability values are large and display little movement under different risk driver positions, such a relative error measure indicates low deviations even if the estimates are quite far off. For example, suppose that the base BEL is at 10,000 and all the actual stresses lead to minimum liability values of 9,950 and maximum values of 10,050. If all the proxy model estimates are off by +100, the relative error measure indicates deviations of a maximum of about 1% even though the proxy model estimates are not very useful, as the proxy model's worst-case BEL stress value is as large as the actual best case BEL. Alternatively, one might consider the relative deviations between BEL(stress) and BEL(base) for the proxy model and the actual values as the error measure. While this approach has its merits, it also has its drawbacks.

The tolerated level of deviation should take into account several important considerations, including:

- Number of risk drivers involved
- Overall number of calibration scenarios
- Magnitude of stress impacts
- Spread of deviations (equally spread or systematically biased)
- Size of deviations in the tail

Taking these aspects into consideration, we concluded that the proxy models generated in the case study pass our quality criteria.

This case study includes dynamic policyholder behaviour as a part of the projection model. Therefore, we see that LSMC is capable of explaining the effects of such dynamic and highly path-dependent actions in a suitable way. Given the nature of dynamic policyholder behaviour or dynamic management actions, this cannot be taken for granted, as these actions are among the most challenging features for the replication of the liability behaviour. To assess the proxy model quality without dynamic lapse, we repeated the case study for both products, this time switching off dynamic lapse in the projection model. The quantitative validations in this case indicate average deviations of 1.2% (GMAB) and 1.9% (GMWB) using the same 37 sample scenarios as above. Hence, we see that dynamic lapse is a factor that makes the calibration of a proper proxy model for the liabilities a more complex task, while on the other hand the loss of explanatory quality is within a reasonable range.

7.4 BUILDING A PROXY MODEL FOR THE SCR

Having calibrated and validated the LSMC proxy models for the GMAB and GMWB products in the previous sections, we now turn to using this model for SCR calculations. The main input here (besides the proxy model itself) is a real-world distribution of the risk drivers. In order to derive the SCR (which, in this case, is the 99.5% quantile of the BEL distribution) we sample a sufficiently large number of joint realisations (in this case, 50,000) of the risk drivers under their real-world probability measure. Each of these realisations is a vector containing the corresponding risk driver sample. Next, the proxy model is evaluated for each of these risk driver vectors and each resulting value constitutes a particular realisation of the liability value. Hence, we get a large number of realisations of the liability value under our real-world view. We then use these to derive the empirical distribution of the liability value and related figures, such as the value-at-risk or expected shortfall for certain probability levels of the SCR. Having a polynomial proxy model available allows for the quick computation of this exercise because all we need to do is to evaluate a polynomial for a large number of risk driver realisations.[7]

[7] Each evaluation of the polynomial corresponds to an assessment of the liability value for the economic situation as determined by the risk driver realisation.

The determination of a real-world distribution for the underlying risk drivers is a complex task for insurance companies. It generally involves taking into account the asset portfolio, calibrating to market conditions, and determining further economic assumptions such as correlation targets. A similar level of complexity applies to the determination of the distribution of non-market risks such as lapse risk. Hence, we acknowledge that there is no unique real-world view and each insurer may have a different one. Taking this into account we perform our case study on a relevant, but not too sophisticated, real-world view to allow for an appropriate illustration of the results. In this case, we assume that the risk drivers follow a multivariate normal distribution with the following parameters:

- Mean values: The mean values of the individual risk factors correspond to the base valuation assumptions.[8]
- Standard deviations: The standard deviations of the individual risk factors have been determined in such a way that the 99.5% quantile of each risk driver coincides with a certain stress target:
  - Interest rates: 99.5% quantile yield curve is initial yield curve stressed with the Solvency II Quantitative Impact Study 5 (QIS5) assumptions.
  - Equity: 99.5% quantile is -39%.[9]
  - Equity volatility: 99.5% quantile is 37.5%, which refers to a multiplicative stress of the base volatility with a factor of 1.5.[10]
  - Lapse: 99.5% quantile is +50%[11] of base case lapses.
  - Bond fund: Change of bond fund is implicitly driven by change of initial yield curves; hence we need no explicit assumptions here.
- By following this approach we calibrate the marginal distribution of each risk factor toward a certain stress target in its tail.
- We used the stylised correlation matrix shown in the table in Figure 26 in our case study[12] in order to describe the dependency structure of the risk drivers.

8 The base valuation assumption for the risk factor PC1 was set equal to -0.5.
9 This is the QIS5 shock level for global equity before application of the symmetric adjustment mechanism.
10 Shock level was chosen for illustrative purposes.
11 This is the QIS5 shock level for lapse. As we assume that the risk drivers follows a normal distribution this is equivalent to using -50% of base case lapses as the 0.5% quantile. We are not making a judgment here as to the more onerous direction of the lapse shock.
12 The distribution of the bond fund implicitly results from the change of the interest rate levels and hence no further assumptions will be required for its distribution.

FIGURE 26: REAL-WORLD CORRELATION TARGETS FOR RISK DRIVERS

	PC1	PC2	Equity return	Equity volatility	Lapse
PC1	1	0	0.5	-0.25	0
PC2	0	1	0.5	-0.25	0
Equity return	0.5	0.5	1	-0.25	0
Equity volatility	-0.25	-0.25	-0.25	1	0
Lapse	0	0	0	0	1

The resulting BEL distributions for the GMAB and GMWB product have the forms shown in Figures 27 and 28.

FIGURE 27: GMAB LIABILITY VALUE DISTRIBUTION

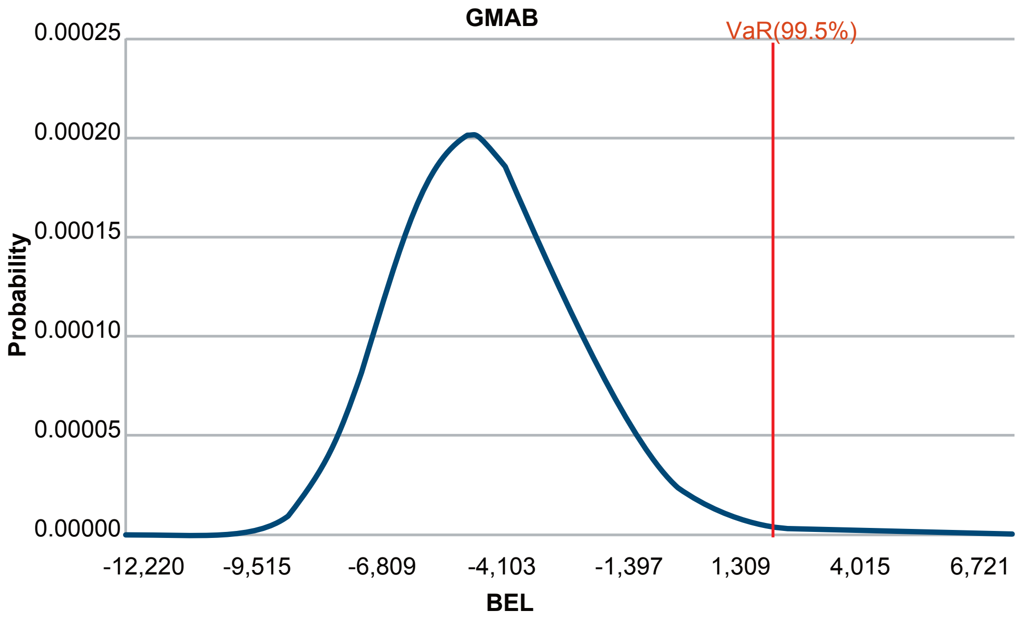

FIGURE 28: GMWB LIABILITY VALUE DISTRIBUTION

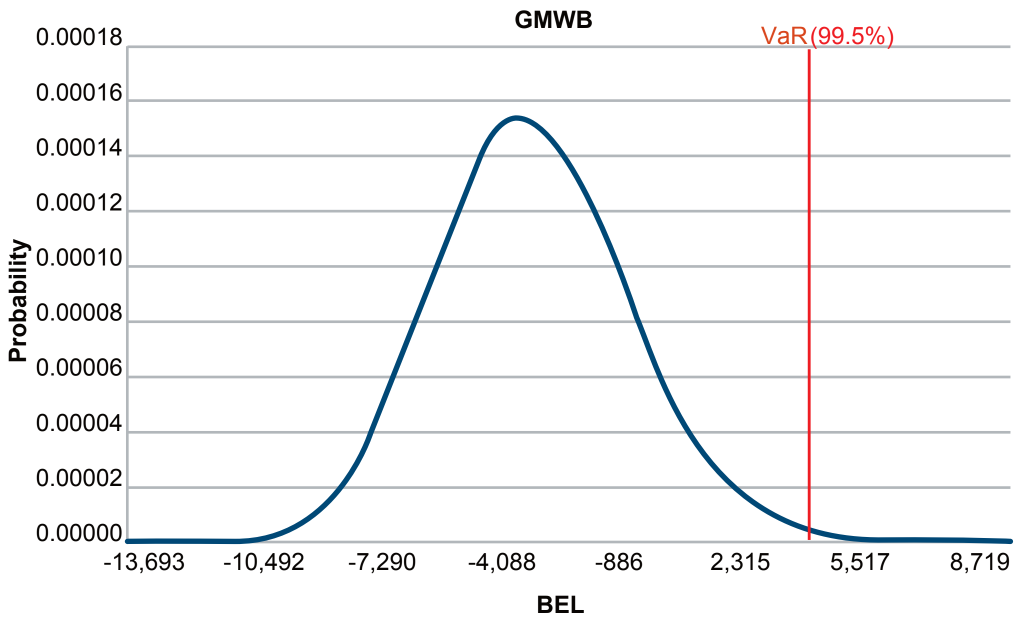

The overall shape of both BEL distributions is right-skewed, indicating that there are asymmetrical rules within the projection models from which the policyholder is benefitting. We see that the level of asymmetry of the BEL distribution is more pronounced for the GMAB product (Figure 27) than for the GMWB one (Figure 28). These distributional shapes are exactly what we expect to see, based on: a) the underlying real-world distribution, and b) the functional shapes as seen in Sections 7.1 and 7.2. The real-world distribution is a multivariate normal one, which has the property that it is symmetrical with respect to the mean value. Hence, any deviation from the mean has equal probability in both directions. Looking at the two-dimensional plots in Sections 7.1 and 7.2, we see that they display either a linear or a convex shape. While a linear function shape evaluated for a set of normal random variables will result in a normal and hence perfectly symmetrical distribution again, a convex shape will result in a right-skewed liability distribution. Convexity refers to the fact that the policyholders benefit more from good scenarios than they suffer from bad ones. Hence, a symmetrically distributed set of risk factors will lead to a right-skewed liability distribution.[13]

From the resulting distributions we get a 99.5% VaR of 1,979 (GMAB) and 4,190 (GMWB). The table in Figure 29 summarises the joint 99.5% change of BEL as well as the marginal 99.5% BEL changes per risk driver. These have been derived by evaluating the LSMC proxy model for corresponding marginal distributions of the risk driver under consideration while all other risk drivers have been kept at their base level. We see that equity return and equity volatility are the most significant risk drivers judged by their individual 99.5% changes of BEL. Taking into account the two-dimensional graphs from the previous section, this is what we expect. The relatively low impact of the bond fund return comes from the fact that even though the bond fund return itself has a significant impact on BEL, its real-world distribution displays quite a low degree of variance.

[13] We can derive this analytically in the following way: If f is a convex function we know that $f(\lambda x + (1-\lambda)y) \leq \lambda f(x) + (1-\lambda)f(y)$ for any λ in [0,1]. By choosing $\lambda=½$ and $y=-x$ this relationship reads $f(0) \leq ½ f(x) + ½ f(-x)$. Without loss of generality we can assume that $f(0)=0$ and get $-f(-x) \leq f(x)$. In the context of our paper, this means that the gain in good scenarios (i.e., the rise of the BEL for large risk driver values *x*) exceeds the loss in bad ones (i.e., the fall of the BEL values for small risk driver values *x*). Thus, if the risk driver X is symmetrically distributed with mean 0, we see that the probability of f(X) exceeding any positive number z is larger than the one of f(X) falling short of -z, which directly leads to a right-skewed distribution of f(X).

FIGURE 29: MARGINAL DISTRIBUTIONS

Total SCR	GMAB	GMWB
BEL	-4,250	-3,170
99.5% VaR	1,979	4,190
Total SCR	6,229	7,360

Breakdown od SCR elements	GMAB	GMWB
Marginal SCR contributions		
Interest Rates (PC1 and PC2)	952	1,725
Equity Return	6,508	7,913
Equity Volatility	2,740	4,879
Lapse	564	57
Bond Fund Return	544	307
Total of Marginal SCRs	11,308	14,881
Diversification Effect	5,079	7,521
Total SCR	6,229	7,360

Note that this is just an illustrative example. Correlations in this example are not intended to represent what might be used in real-life situations.

A straightforward validation of the SCR itself would involve a brute-force nested stochastic approach. Avoiding such an approach is the main reason for the application of proxy models and hence this form of validation is unsuitable. Therefore, it is essential to have a proper validation of the proxy model itself before turning to an assessment of the SCR. Once we're comfortable with the results of the proxy model validation, i.e., once we have seen that it passes validation testing for the change of liability value under a variety of economic conditions (as seen in Figures 10 to 25) with a reasonably low level of variance (as assessed by the confidence intervals in Figures 16,17, 23, and 24), we can apply the proxy model to estimate the distribution of the liability value under a particular real-world distribution of risk drivers. Note that there are methods which may be used to explicitly assess the variance of the SCR estimator (see Kalberer in the References section for details) but we do not cover them further here.

Alternative joint real-world distributions of the risk drivers are typically either:

- Copulas to join assumed[14] marginal distributions of risk drivers[15]
- Causal model structures where a causal structure of the risk drivers and the underlying factors driving those risks are modelled jointly and simulated by a set of underlying assumptions

14 For example, distributions per risk driver that are calibrated to a certain set of parameters and assumptions for the underlying asset portfolio.

15 Note that the approach of using a multivariate normal distribution as a real-world distribution is equivalent to joining normally distributed marginal distributions using a Gaussian copula.

7.5 INCLUSION OF FURTHER RISK DRIVERS OUTSIDE OF THE MODEL

Typically, some risks such as operational risk are not part of the actuarial cash-flow model and hence cannot be directly included in the LSMC fitting and calibration process. Using some simplifying assumptions, such risks can be aggregated with those covered by the cash-flow model and/or LSMC proxy model:

- For example, assume that:
  - Any such risk from outside the model has a linear impact on the liability value
  - The SCR values for the individual 'outside risks' have been aggregated outside the model – to a value SCR_{out}
  - The overall impact of the risks from outside the model is normally distributed and SCR_{out} displays correlation ρ with SCR_{LSMC}
- Calculate the overall SCR_{LSMC} as done above excluding 'outside risks' and aggregate both elements of the SCR as follows: $SCR_{total}=\sqrt{SCR^2_{LSMC}+SCR^2_{out}+2\rho SCR_{LSMC}SCR_{out}}$.

Even though the assumptions above will typically not hold from a theoretical point of view, they may be regarded as a reasonable proxy of the true circumstances and applied, taking into account that there might be no perfect solution for this problem.

Other approaches can be used to join the marginal distribution of the liability value under the risks modelled via LSMC and risks that are outside the proxy model. Basically, this comes back to the problem of joining different marginal distributions (where at least one of them–the distribution of the liability value under those risks which are part of the proxy model–is non-analytic) and can be addressed with similar approaches, as discussed in the previous section (dealing with the joining of real-world risk driver simulations).

It is useful to consider an example of the possible form of a causal model for joining an outside risk with the liability value distribution covered by the proxy model. Suppose Company A holds equity as well as a bond issued by another company, Company B. For simplicity, we assume that the bond debt will be repaid by Company B after one year. Company A considers this repayment subject to default risk and assumes that in the case of default the liability value will increase by 500.[16] The default risk is not part of the cash-flow model of Company A and hence cannot be captured by the LSMC proxy model. However, Company A has a separate 'asset-only' model to simulate the defaults of Company B. This model follows a Merton-like approach for credit risk (see Merton in the References section for details) and simulates the assets of Company B under a 1one-year real-world view in a very simplistic way (by considering the assets of Company B as a single equity position, with a drop in this equity position below a certain threshold indicating the default of Company B).

Since Company A also holds an equity stake in Company B, it models its real-world view on the overall equity return (which is a risk driver of the proxy model) by joining the real-world performances of the shares of Company B and a global equity index (which refers to the performance of its remaining equity positions).

16 Of course, this is a simplification of reality where the default of an individual bond will have an impact on liabilities, which depends on other risk drivers as well.

Hence, Company A has generated 10,000 realisations of the BEL under the risks that are part of the proxy model (see first block of the table in Figure 30 below) including EquityB, the performance of Company B's equity that forms part of the overall equity performance. It uses these 10,000 realisations of EquityB to derive 10,000 simulations of the default events for the bond issued by Company B (see second block of Figure 30). It then joins the influence of the default of Company B on the overall BEL by adding the change of BEL caused by the default of Company B to the BEL derived by the proxy model (see third block of Figure 30). Therefore, the insurance company has joined the distribution of the BEL under the risks covered by the proxy model and the change of BEL distribution under an outside risk in a consistent way by linking the performances of the equity of Company B that (partly) drives the overall equity risk as well as the default of Company B's debt.

FIGURE 30: SCENARIO AGGREGATION

Simulation	EquityB	$\text{Equity}_{\text{global}}$	…	BEL	EquityB	Default	Change of BEL	Overall BEL
1	-13%	+4%	…	10,000	-13%	No	-	10,000
2	-45%	+4%	…	9,500	-45%	Yes	+500	10,000
3	+12%	+15%	…	7,000	+12%	No	-	7,000
4	+0%	-6%	…	11,000	+0%	No		11,000
…	…	…	…	…	…	…	…	…
9,999	+4%	+7%	…	8,000	+4%	No		8,000
10,000	-29%	-3%	…	10,500	-29%	Yes	+500	11,000

8 LSMC ACROSS A MULTI-YEAR TIME HORIZON

8.1 APPLICATIONS

In the previous section, we discussed how to apply LSMC to deal with one-year VaR problems–in other words, the time horizon had been t=1. While those problems are at the core of many calculations, such as those required for Solvency II capital requirements, several instances where proxy models may be applied require consideration of longer time horizons. In this section we briefly discuss some examples of such applications.

8.1.1 OWN RISK AND SOLVENCY ASSESSMENT

According to Article 45 of the Solvency II Directive (see Solvency II Directive in the References section), every insurance undertaking shall carry out an ORSA. In particular:

…The own risk and solvency assessment shall be an integral part of the business strategy and shall be taken into account on an ongoing basis in the strategic decisions of the undertaking.

Furthermore, Consultation Paper 008/2011 issued by the European Insurance and Occupational Pensions Authority (EIOPA) provides the following interpretation of the directive:

…The undertaking's assessment of the overall solvency needs should be forward-looking.

…The undertaking needs to project its capital needs over its business planning period. This projection is to be made considering likely changes to the risk profile and business strategy over the projection period and sensitivity to the assumptions used.

In other words, Solvency II coverage ratios should be assessed over a period extending beyond one year (perhaps three to five years). There are a number of ways in which such projections may be prepared, such as the projection of balance sheets and capital requirements along certain deterministic paths that are considered appropriate for business planning purposes.

8.1.2. VALUE-BASED MANAGEMENT

Valuation metrics such as market-consistent embedded value (MCEV) represent a valuation exercise with time horizon t=0, while Solvency II SCR calculations at the valuation date involve a one-year time horizon. However, the availability of MCEV and SCR numbers is generally not enough for the purposes of value-based management (VBM).

Naturally, one can make use of information available from sensitivities and analysis of change for VBM purposes. However, these sources do not provide further information such as how MCEV or other economic balance sheet items would evolve over a multi-year planning period. This is where a proxy model can prove useful.

8.2 CHALLENGES OF MULTI-YEAR CALCULATIONS

The problems mentioned above imply a number of practical challenges, which we will briefly touch on in this section. A thorough treatment of these challenges would go beyond the scope of this current paper.

8.2.1. MODELLING FUTURE NEW BUSINESS

Usually, valuation exercises at t=0 are concerned with the business in-force–for example, MCEV principles clearly exclude future new business from the MCEV scope. However, a meaningful consideration of MCEV development over a multi-year business planning period cannot be carried out under the assumption that the insurance undertaking will not write any new business during that period. The same reasoning applies to the question as to how the main risks, measured by their solvency capital requirements, would evolve over such a medium-term time horizon.

Modelling future new business leads to a number of interesting questions, such as how to set assumptions for future new business (for example volume and profitability) or expanding upon dynamic management actions in the cash-flow model in order for the latter to cover such actions as a reduction of new business volume in 'bad' scenarios.

8.2.2. LSMC PROXY MODELLING AT A TIME HORIZON $t > 1$

Generally speaking, the LSMC process lends itself to an arbitrary time horizon very well. In particular, it is possible to obtain an adequate liability function for any relevant time horizon. However, the $t > 1$ case includes an additional subtlety which we did not have to worry about when dealing with the 1-year time horizon. The economic situation of an insurance undertaking at time horizon $t > 1$ depends not only on the current risk driver values but also on the history between the valuation date and time *t*–whereas in the case of $t = 1$, all scenarios for $t = 1$ share the same history at $t = 0$. For example, it is not enough to know that the interest rates are 5% flat in some scenario for $t = 3$ without also knowing how high the interest rates have been at times $t = 1$ and $t = 2$ or whether, for example, an insurance undertaking was able to declare a bonus to with-profits policyholders at times $t = 1$ and $t = 2$, and so on.

Thus, we have to introduce additional risk drivers–called state variables–in order to ensure that the LSMC process can properly distinguish between two scenarios which may be similar or even identical at time horizon t and yet feature very different history between time 0 and time t.

We cannot determine universally appropriate state variables, because their choice is clearly dependent on the particular problem and the underlying stochastic cash-flow model.

There are certain key questions and considerations to bear in mind when preparing the model for multi-year calculations:

- Identification of important state variables: Which variables at time t most efficiently incorporate the history of a certain scenario prior to time t?

- Assessment of the correlation of the state variables: Is all of the information contained in the state variables necessary or is there a certain level of redundancy when using all of them? If the state variables display high correlations amongst each other, then it may be possible to remove certain variables that are highly correlated with others in order to make the analysis more manageable without significant loss of accuracy.

- Check that the future risk factor distributions within calibration scenarios are rich enough: Just as for the case at time $t=1$ we need a broad universe of risk factor positions to be covered over the projection period. This is necessary to ensure that we can explain the liability value during the projection at all different risk factor levels and are able to stress those risk factors to assess the capital requirements at any point in time and for any scenario.

For a more thorough discussion of this topic, please see Bettels et al. in the References section below.

8.2.3. ESTIMATION OF SOLVENCY II COVERAGE RATIOS AT t > 1 FROM LSMC RESULTS

In order to obtain a Solvency II coverage ratio for a particular scenario s and time horizon t, one has to define an approach as to how to calculate this coverage ratio based on some LSMC liability functions such as BEL or present value of future profits (PVFP), which can be evaluated for that scenario and time horizon.

In other words, one has to define an approach as to how to calculate *SCR(t,s)* as well as the available capital *AC(t,s)*.

It is possible to calculate *SCR(t,s)* using a pragmatic approach such as the following:

- Evaluate the LSMC liability function at time *t* in order to obtain the base value *PVFP(t, s)*. Here, the state variables would take the history of scenario *s* between time 0 and time t into account.

- Evaluate the LSMC liability function for time *t* for each SCR-relevant stress in order to obtain an individual *SCR(i)* for each risk i. The magnitude of each stress would represent the corresponding ‘one-in-200-year’ event for each risk.

- Aggregate the individual *SCR(i)* obtained in the previous step via a standard formula-style correlation approach.

Alternatively, one could construct many thousands of real-world scenarios emanating from (t,s) over the one-year horizon, evaluate the LSMC liability function on each of them, and calculate SCR(t,s) from the resulting risk distribution. However, this approach is likely to be extremely burdensome in practice.

In order to calculate the available capital *AC(t,s)*, we would need the value *PVFP(t,s)* calculated above and the risk margin *RM(t,s)*. A pragmatic way of estimating the latter would be a cost of capital approach. Here, we obtain $SCR^{nh}(t,s)$ for non-hedgeable risks as outlined above. We then compute $SCR^{nh}(T,s)$ for T > t by projecting $SCR^{nh}(t,s)$ along a simple risk driver such as value of reserves in the best estimate scenario (i.e., a scenario in which future interest rate outcomes match exactly the forward rates predicted by the yield curve at the outset).

9 SOLVENCY I CAPITAL

Owing to the recent delays to the planned implementation of Solvency II regulations, it is likely that companies will continue to need to be able to calculate their capital requirements under Solvency I rules for some time.

9.1 VARIABLE ANNUITY CAPITAL REQUIREMENTS

Currently the application of Solvency I rules throughout Europe varies to a certain degree by territory. In addition, the current rules were introduced before the recent development of the VA market in Europe and so do not capture particularly well the types of risks inherent in VA products.

In recent years, some individual regulatory authorities have introduced specific capital requirements to deal with the exceptional features of VA products. In Ireland, which is one of the primary jurisdictions from which VA business is written in Europe, the regulator (the Central Bank of Ireland) requires companies to hold total capital (reserve plus solvency margin) on new business written to cover what is known as the 'CTE95' outcome. 'CTE' stands for conditional tail expectation, while '95' indicates that it is the average of the worst 5% of outcomes that is of interest. The methodology is a real-world stochastic methodology, involving the projection of future income and outgo over a range of different scenarios. Economic assumptions are based on estimates of future market conditions while other assumptions (such as demographic experience) are based on prudent estimates adjusted to capture any specific requirements of the regulations.

The Irish approach has been adapted from the US approach. The umbrella body for insurance supervisory authorities in the US, the National Association of Insurance Commissioners (NAIC), asked the American Academy of Actuaries (AAA) to produce standards for setting the level of capital required to support VA liabilities. Over the course of a number of years the AAA has formulated and refined an approach to setting VA capital requirements.

At a high level, the US methodology involves running a cash-flow model over a large number of investment scenarios (typically 10,000), calculating a present value of future losses for each scenario and basing the total capital requirement on the distribution of those losses. More specifically, the methodology involves projecting policy cash flows stochastically across a large number of scenarios and then determining the greatest present value of the cost of the guarantees for each scenario. The total capital requirement is set to the average of the 10% of scenarios which produce the largest negative present values—which is generally referred to as the CTE90 (conditional tail expectation of the 10% of scenarios which produce the largest negative present values). The assumptions underlying these scenarios, for example the mean paths and volatilities of the projected assets, are external parameters and fixed for each projection. Therefore, one can characterise this situation as follows: The liability value is the CTE90 emerging under a real-world distribution where the parameters of this distribution serve as risk drivers. Note that this kind of liability value is quite different to an economic value in the Solvency II context. However, this definition is suitable for the purposes of applying a proxy model, as we outline below.

9.2 PROJECTION OF SOLVENCY I CAPITAL REQUIREMENTS

An LSMC approach to the projection of capital requirements in a Solvency I context will look similar to the following:

- Risk drivers: Key parameters of the underlying real-world distribution used to project the VA cash flows serve as risk drivers. Examples are:

  — Initial yield curves
  — Mean paths of interest rates and indices
  — Volatilities of interest rates and indices

- Outer scenarios: Just as for the Solvency II type applications in Section 5, we identify a reasonable range of values for each risk driver. The *n* outer positions will then be uniformly distributed over the multidimensional space of risk drivers, i.e., multidimensional parameter vectors that span the space of values under consideration as uniformly as possible.

- Inner scenarios: For each outer scenario, i.e., for each combination of parameters for the underlying real-world distribution, we generate a set of *Y* inner scenarios which incorporate the corresponding outer assumptions.

- Valuation: For each of the *n* outer positions we derive the present value of future losses for each of the Y inner scenarios and derive the CTE90 by taking the average of the 10% of scenarios which produce the largest negative present values of future losses. Any such value will serve as CTE90 estimate per outer scenario.

- This puts us in the same position as discussed in Section 5.1—for each outer scenario, i.e., each parameter combination of the real-world scenario distribution, we get a rough liability value estimate and use the LSMC approach to smooth out the noise and uncover the true relationship between these parameters and the resulting CTE90. Note that the LSMC approach does not rely on the fact that the liability value is related to the valuation of liability cash flows under a risk-neutral probability measure. All it takes is a liability value that is derived as a (conditional) mean over a set of stochastic scenarios.

Note on convergence: As seen in Section 5.2.4, the underlying assumption of LSMC was that the error term, i.e., the difference between the actual liability value and its estimate derived by using Y inner simulation, is normally distributed. For valuations in the context of Solvency II, we can assume this assumption to be approximately fulfilled when using even very few inner scenarios. Empirically, this means that the estimated liabilities encountered over the range of outer scenarios will tend to be too large with the same probability as they tend to be too small, and hence this symmetrical error can be cancelled out through regression.

However, in the context of Solvency I, where the liability value is derived from the 10% worst-case inner scenarios, this will require a significantly larger number of inner scenarios per outer one because the estimation of the CTE90 based on only a small number of inner scenarios will display a significant bias toward smaller values. This comes from the fact that, when estimating the CTE90 with only a few scenarios, we are unlikely to encounter the really large losses that can have a significant influence on the conditional tail expectation, and thus the resulting estimated conditional tail expectations tend to be too small. Therefore, the error is non-symmetrical, which introduces a bias to the resulting LSMC proxy model. It is advisable to perform a thorough analysis of the empirical distribution of the error, ε, by considering the realisations of the error variable (also referred to as residuals) $\varepsilon_i = Y^{(i)} - f(X_1^{(i)},\ldots,X_m^{(i)})$, $i=1,\ldots,n$. To counter this effect we have to use significantly more inner scenarios per outer one and use the distribution of the errors in order to apply a (subjective) assessment of whether or not the chosen number of inner scenarios is sufficient to overcome the estimation bias described above.

10 REFERENCES AND CONTACT DETAILS

10.1 REFERENCES

Ledlie, C., Corry, D., Finkelstein, G., Ritchie, A., Su, K., & Wilson, D. (2008). Variable Annuities. Institute of Actuaries paper.

Bettels, C., Hörig, M., Ketterer, F., & Leitschkis, M. (2013). Using least-squares Monte Carlo in a multi-year context. *InsuranceERM*.

Hörig, M., & Wechsung, M. (2013). Scenario Reweighting Techniques for the Quick Re-Calibration of Pricing Scenarios: Applications in Insurance Risk Management.

Jolliffe, I.T. (2002). Principal Component Analysis. New York: Springer.

Kalberer, T. (2012). Stochastic determination of the value at risk for a portfolio of assets and liabilities (part 1, 2 and 3). *Der Aktuar*, 01/02/03.2012.

Merton, R.C. (1974). On the pricing of corporate debt: The risk structure of interest rates. *Journal of Finance*, 29.

Solvency II Directive 2009/138/EC.

Wu C.F.J. (1986). Jackknife, Bootstrap and other Resampling Methods in Regression Analysis, Annals of Statistics 14, pp. 1261-1350.

10.2 CONTACT DETAILS

The contact details for the authors and reviewer of this report are as follows:

Mario Hörig, Consultant, Düsseldorf
mario.hoerig@milliman.com
+49 211 9388 6613

Karl Murray, Consulting Actuary, Dublin
karl.murray@milliman.com
+353 1 647 5509

Eamonn Phelan, Principal, Dublin
eamonn.phelan@milliman.com
+353 1 647 5914

Michael Leitschkis, Practice Leader, Düsseldorf
michael.leitschkis@milliman.com
+49 211 9388 6618

Casey Malone, Consulting Actuary, Seattle
casey.malone@milliman.com
+1 206 504 5921

APPENDIX

A.1. DETAILS OF THE PRODUCTS USED IN THE CASE STUDY

Note that all of the assumptions below are simply for illustrative purposes and do not purport to be representative of current or expected future pricing conditions.

A.1.1 GMAB

The guarantee is a return of premium on death or after 10 years.

Guarantee charge = 1.4% p.a. (levied on fund value).

The initial split of the fund is 35% equities and 65% bonds.

Single premium = €100,000.

Age at entry = 50.

Fund management charge = 1.6% p.a.

The maintenance expense = €40 p.a. inflating at 2.5% p.a.

Lapses:

Base rates (p.a.)

Year	GMAB
1	0.0%
2	4.5%
3	6.0%
4	6.0%
5+	7.0%

Dynamic behaviour:

Effective lapses = max (Base lapses * Dynamic lapse factor, 2%)

Where

ITM = Account Value / Benefit Base

Dynamic lapse factor = min (140%, max (0%, 2 * ITM – 1))

FIGURE 31: GMAB DYNAMIC LAPSE FACTOR

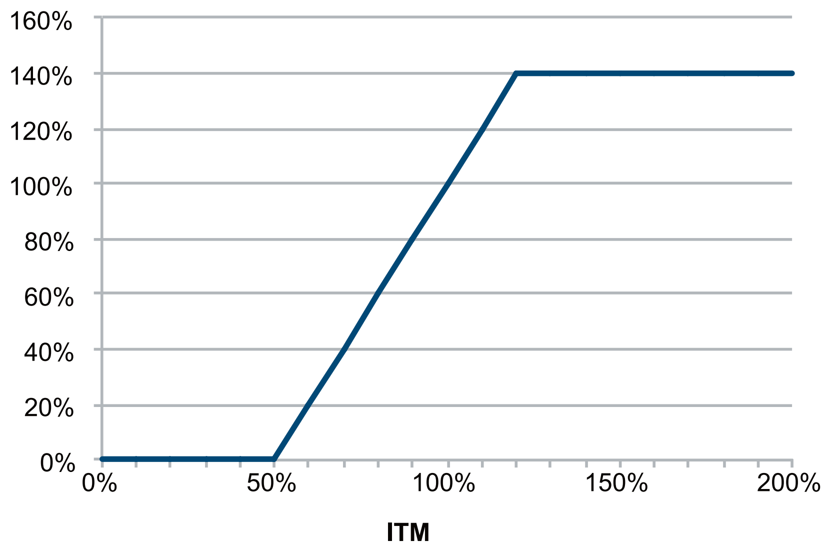

A.1.2 GMWB

Details of the guarantee are as follows:

- There is a deferral period of five years and withdrawals start after this period
- During the deferral period the benefit base rolls up at 2.5%
- An annual ratchet applies
- The guaranteed withdrawals after the deferral period are 3.5% per annum of the benefit base

Guarantee charge = 1% p.a. (levied on fund value).

The initial split of the fund is 50% equities and 50% bonds.

Single premium = €100,000.

Age at entry = 60.

Fund management charge = 1.6% p.a.

The maintenance expense = €40 p.a. inflating at 2.5% p.a.

Lapses:

Base rates (p.a.)

Year	GMAB
1	1.0%
2	2.0%
3	3.0%
4	4.0%
5+	5.0%

Dynamic behaviour:

Effective lapses = Base lapses * Dynamic lapse factor

Where

ITM = 10 year spot rate – Withdrawal Rate * Benefit Base / Account Value

Dynamic lapse factor = min (200%, max (0%, (100 / 3) * (ITM + 2%)))

FIGURE 32: GMWB DYNAMIC LAPSE FACTOR

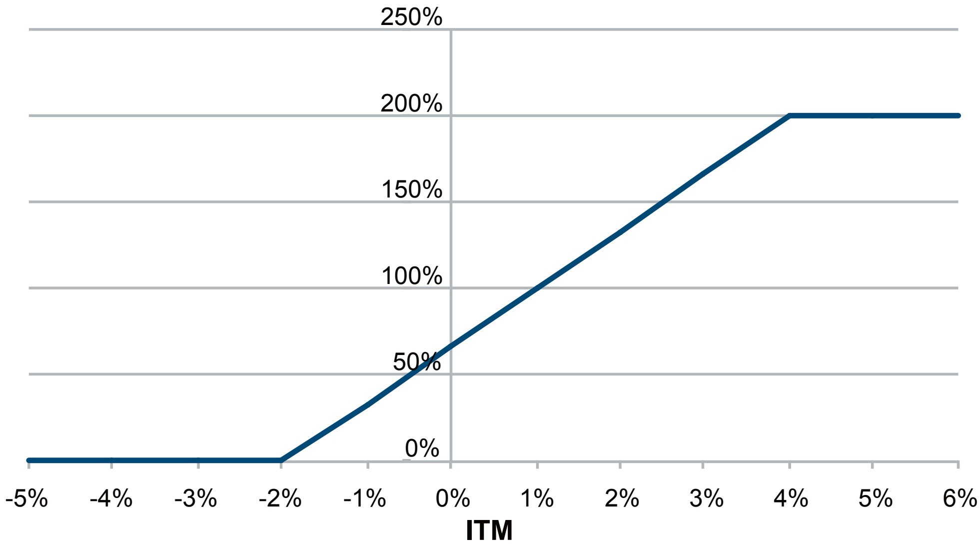

A.2. ECONOMIC ASSUMPTIONS

Lognormal equity returns with 25% volatility

Single factor Hull-White model for interest rates with mean reversion 3% and volatility 1%[17]

0% equity/interest rate correlation

Bonds modeled as 7-year risk-free zero coupon bonds

The yield curve that was used for the pricing of the model points used in the case study represents the Q3 2012 yield curve (the 'market curve' in the chart in Figure 33).

For the purposes of calibration of the LSMC proxy model, a higher yield curve ('PC1=0' in Figure 33) was used in order to avoid problems associated with negative interest rates, which would arise if the market yield curve had been used for this purpose.

[17] J. Hull, A. White. (1990). Pricing Interest Rate Derivative Securities. *Review of Financial Studies*.

Figure 33 shows that PC1=-0.5 is a good approximation for the market curve.

FIGURE 33: DISCOUNT RATE CURVES

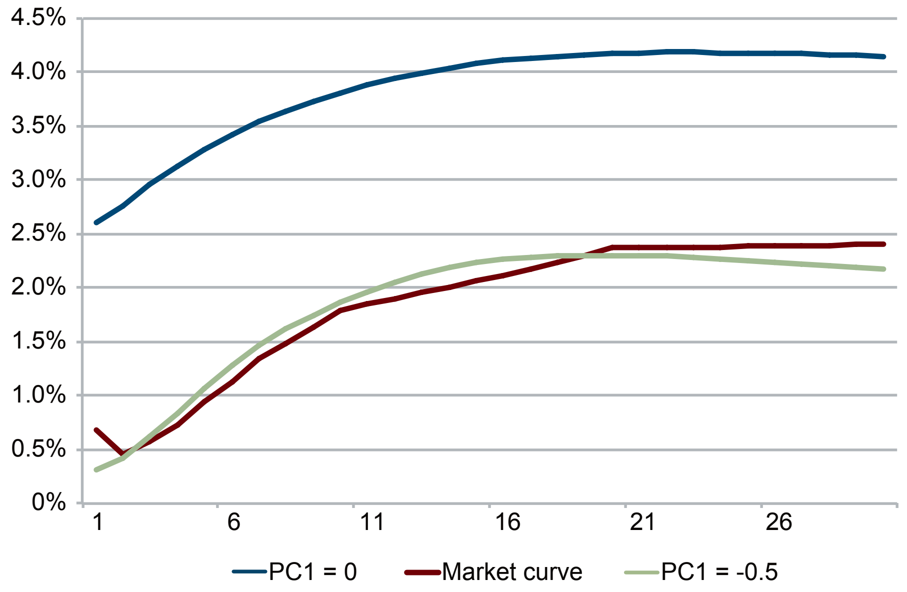

A.3. BACKGROUND ON JACKKNIFING THE LSMC CONFIDENCE INTERVALS

Suppose we've derived a liability function *f*. For a certain risk driver combination $X_1,\ldots,X_m$ we get an estimate for the corresponding liability value by evaluating $f(X_1,\ldots,X_m)$ (called a *point estimate*). In order to derive confidence intervals for the point estimate $f(X_1,\ldots,X_m)$ we apply a certain resampling procedure, called jackknifing. The jackknifing aims at subsequently resampling the liability function from the simulation data used for the LSMC fitting. Each resampling results in a new liability function and a new point estimate that is used to derive standard deviations and confidence intervals for the point estimates of the liability function. The algorithm for the jackknifing is as follows (see Wu):

Input

- *n* liability values $Y^{(1)},\ldots, Y^{(n)}$ and the values of the corresponding risk driver vector $X^{(i)}=(X1^{(i)},\ldots,Xm(i))$, i=1,…,m
- The final polynomial f from the LSMC fitting approach (i.e., a vector of coefficients $\beta=(\beta_1,\ldots, \beta_k)$ and corresponding polynomial)
- Number *p* of jackknife repetitions (e.g., p = 5,000)

Process

Repeat the following steps *p* times (with the current step indicated by *j*):

- Draw a random subset of the fitting data with size d=floor(0, 8*n) (i.e., the subset consists of floor(0, 8*n) distinct liability values from [$Y^{(1)}$,…, $Y^{(n)}$] and their corresponding risk drivers, where floor[x] refers to the largest integer smaller or equal to *x*). In this case 'random' means that each liability value may be drawn with the same probability.

- Refit the final polynomial structure from the original LSMC fitting to this data.

  – This results in a new vector of coefficients $\hat{\beta}^{(j)}$.

  – Determining $\hat{\beta}^{(j)}$ involves inverting the matrix $X_j^T X_j$, where X_j is built in the same way as the final design matrix *X* (i.e., the matrix *X* used for the final fit that determines β) but consists only of *d* of the *n* lines of *X*. Calculate $v_i = \det(X_j^T X_j)$, the determinant of $X_j^T X_j$.

  – v_j is the specific weight of results from j^{th} jackknife trial.

- Compute $\beta^{(j)} = \hat{\beta}^{(j)} + \sqrt{\left(\frac{d\text{-}k+1}{n\text{-}d}\right)}\left(\hat{\beta}^{(j)} - \beta\right)$, where the operations are applied component-wise.

  – The vector $\beta^{(i)}$ represents the liability function estimated in the current jackknife trial.

Output

This procedure yields p different liability functions $f_1,\dots,f_p$ (represented by the corresponding coefficient vectors $\beta^{(1)},\dots,\beta^{(p)}$) and their associated weights $v_1,\dots,v_p$.

This information can be used to derive confidence intervals and standard deviations for $f(X_1,\dots,X_m)$ for any arbitrary risk drivers value $X=(X_1,\dots,X_m)$. In order to do so we proceed in the following way:

- Calculate the values $f_1(X_1,\dots,X_m),\dots,f_p(X_1,\dots,X_m)$ by using the polynomial structure of *f* and the coefficient vectors $\beta^{(1)},\dots,\beta^{(p)}$.

- Set $w_i = \frac{v_i}{\sum_{j=1}^{p} v_j}$ (i=1 ,..., p) → individual weight of the i^{th} jackknife trial.

- Calculate standard deviation $[f(X_1,\dots,X_m)] = \sqrt{\sum_{j=1}^{p} w_j \left(f\left(X_1,\dots,X_m\right) - f_j\left(X_1,\dots,X_m\right)\right)^2}$.

- To sample a confidence interval of $f(X_1,\dots,X_m)$ at any confidence level α, we start by sorting the values $f_1(X_1,\dots,X_m),\dots,f_p(X_1,\dots,X_m)$ in ascending order and considering their associated weights. The α confidence level ($0 \le \alpha \le 1$) is derived by choosing the smallest function value such that the sum of its weight and the weights of all smaller function values exceed α.

Milliman

Mario Hörig
mario.horig@milliman.com

Karl Murray
karl.murray@milliman.com

Eamonn Phelan
eamonn.phelan@milliman.com

Michael Leitschkis
michael.leitschkis@milliman.com