MOODY'S ANALYTICS

RESEARCH
OCTOBER 2013

Proxy Function Fitting: Some Implementation Topics

Gavin Conn FFA
Moody's Analytics Research

Contact Us

Americas
+1.212.553.1658
clientservices@moodys.com

Europe
+44.20.7772.5454
clientservices.emea@moodys.com

Asia (Excluding Japan)
+85 2 2916 1121
clientservices.asia@moodys.com

Japan
+81 3 5408 4100
clientservices.japan@moodys.com

Executive Summary

Liability proxy function fitting - typically using Least Squares Monte Carlo, curve fitting or some mix of the two – is increasingly a business-as-usual activity in economic capital modeling across the global insurance sector. Much progress has been made in recent years in developing the theoretical framework for efficient proxy function fitting, and in the development of software products that can reliably implement and automate the required workflows.

During our significant experience in supporting clients in the development of their liability proxy functions, there are some general technical themes that often arise in the practical implementation. This note provides discussion and examples in some of those areas.

Table of Contents

1. Introduction

Proxy functions are increasingly used by insurers in their economic capital models as they enable the insurer to quickly re-value their liabilities under a range of scenarios. Two common approaches used to produce proxy functions are curve fitting and Least Squares Monte Carlo (LSMC). This paper explores some of the practical issues that an insurer may encounter when implementing proxy functions. This discussion is developed through the use of a case study.

Section 2 introduces the case study; sections 3 and 4 discuss two practical implementation topics: modelling demographic risks and yield curve mapping respectively; and section 5 concludes.

2. Case Study

The case study involves a pension product that has a guaranteed annuity option that is applicable to the fund value at a future maturity date. We assume 10 years until maturity and the initial value of the fund is £100,000. The fund is invested in a mix of GBP equities and government bonds, with an equity allocation of 40%. At maturity a guaranteed annuity option applies with a guaranteed annuity rate of 10%, i.e. the policyholder has the option to convert their fund in to an annuity at a rate of £1 per annum for every £10 of fund[1]. Furthermore, the policy carries an additional "money-back" guarantee which ensures that the policyholder will receive a minimum of £100,000 at maturity (this minimum fund level also applies if the GAO is exercised). A deterministic lapse assumption of 5% per annum is assumed in the model, and we ignore the effect of mortality prior to maturity.

The following market risk variables are considered:

- *Equity risk*: a single GBP equity index is modeled. The dependency of the liability value to equity risk is complex. In a low yield environment the GAO is in-the-money and high equity returns increases the fund to which the GAO applies, thereby increasing the shareholder cost. In high yield environments when the GAO is out-of-the-money, it is low equity returns which increase the shareholder cost as the minimum maturity guarantee will bite.
- *Nominal interest rates*: two risk factors are modeled for interest rates. Broadly speaking the first factor controls the short rate and the second factor the long rate. Given that annuity rates are dependent on long-term government bond yields, it is the second factor which plays a more prominent role in the liability value. There is a negative relationship between the liability value and the second interest rate factor – high values imply high yields which reduce the GAO cost.

The B&H Economic Scenario Generator (ESG) is used to value the liability. For this market-consistent valuation nominal interest rates are modeled using the 2-factor Libor Market Model (LMM) and equities are modeled using the Time Varying Deterministic Volatility (TVDV) model. At 31/12/12 using the B&H standard calibration the market-consistent valuation of the guarantee is estimated at £23,800. Note this guarantee cost has two underlying components: the money-back fund value maturity guarantee, and the annuity guarantee. The high guarantee cost is primarily driven by the deep in-the-money guaranteed annuity rate of 10% – the at-the-money guaranteed annuity rate was only 6.25% at the valuation date (which implies a pension which is 37.5% lower than that provided under the guarantee).

To evaluate the liability proxy function the technique of Least Squares Monte Carlo (LSMC) is applied using the B&H Proxy Generator software. A description of the LSMC technique and its application to insurance liabilities is covered in detail elsewhere[2] and therefore only a high level description is given below. In the context of our case study the steps involved in the process are:

1. *Selection of the risk drivers*. Initially we restrict our modeling solely to market risks and consider three market risk drivers, which comprise a single equity index and two factors representing the GBP nominal yield curve. The liability value will of course be a function of a larger number of variables but, for the purpose of illustration, we will only model three market risk factors. The example could be easily extended to cover other risks such as interest rate volatility and equity implied volatility. In this example the dimensionality of the yield curve has been reduced to two underlying risk factors – given that the yield curve valuation model is a two factor stochastic model, it is convenient to use the two underlying risk factors as arguments in the proxy for the liability function.
2. *Creation of the fitting scenarios*. Fitting scenarios are used to create approximate liability valuations which will be used in fitting the proxy formula to the true liability value. In creating fitting scenarios, we choose a multi-dimensional range over which we would like to approximate the true liability function and then create a set of uniformly distributed fitting scenarios. Each scenario contains a random change to the value of each of the chosen liability risk drivers over a 1 year horizon. A large

[1] The policyholder is assumed to have a life expectancy of 20 years at maturity of the fund.

[2] Refer to the B&H research paper 'A Least Squares Monte Carlo Approach to Liability Proxy Modelling and Capital Calculation' *(Koursaris, 2011)*

number of these stressed positions is created to obtain the *outer fitting scenarios*. Once we have created these stressed fitting positions, we re-calibrate our risk neutral economic scenario generator using the stressed risk drivers e.g. using the new yield curve and equity value. We then run a small number of random scenarios of our risk neutral economic scenario generator for each of the fitting positions, to produce the *inner fitting scenarios*. In this case study we have used 20,000 outer scenarios, each with two inner scenarios.

3. *Liability valuation*: The liability is evaluated for each outer fitting stress. The shareholder guarantee cost is evaluated as the average discounted shortfall across the two simulations[3], taking in to account both the minimum guaranteed fund value and the guaranteed annuity option.

   The graphs below show the liability values for each outer fitting stress plotted against the risk driver value. The strong relationship between the liability value and the second interest rate factor is apparent.

   Figure 1: Valuations for the fitting stresses

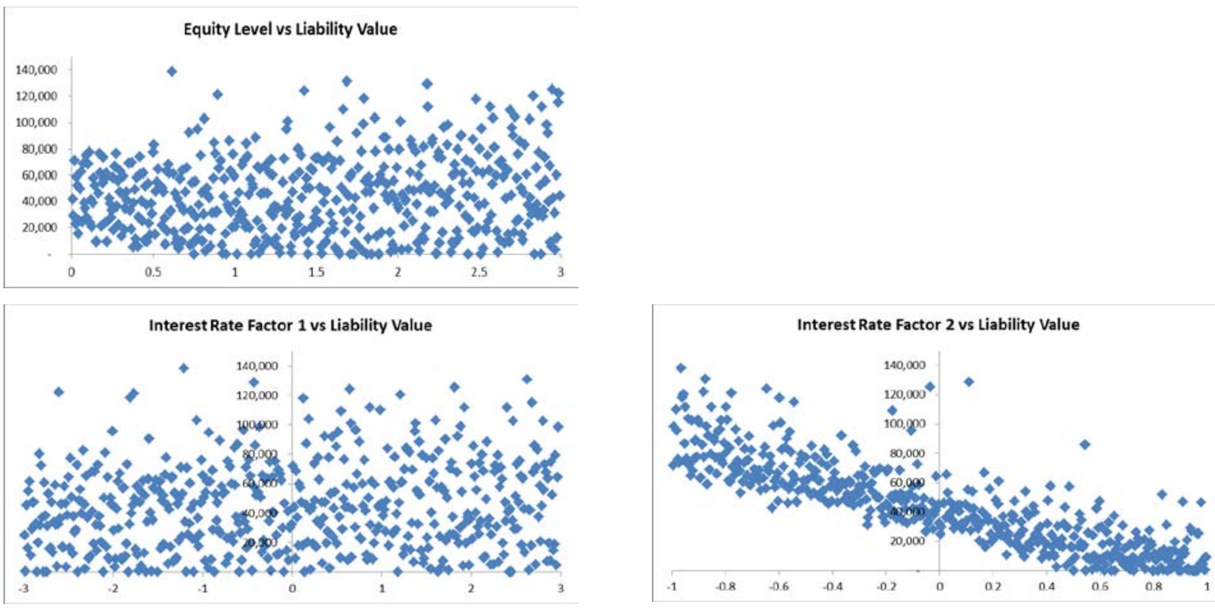

4. *Regression across the liability values to obtain the proxy function*. The proxy generator uses a regression algorithm to construct a liability proxy function that captures how the liability value reacts to changes in the underlying risk drivers. The optimally fitting polynomial function of risk drivers is constructed from the set of allowable polynomial terms. The proxy is restricted to be a polynomial with a maximum of 25 terms and to terms of order four.

   Guarantee Cost = *f (Equity level, Interest Rates factor 1, Interest Rates factor 2)*

Validating the Proxy Function

To validate the proxy function we look at 20 validation scenarios which cover both univariate and multivariate stresses to equity level and nominal interest rates. The equity level stresses range from -50% to +50%, and the yield curve stresses range from a parallel stress of (100)bp to +350bp, as well as twists to the yield curve. The validation scenarios involve a "full blown" revaluation of the liability using Monte Carlo simulation with 5,000 simulations.

Figure 2 shows an excellent fit—in most of the validation scenarios the proxy value is within 2% of the actual value. In this initial example we have only looked at three market risk drivers and the validation implies that 20,000 fitting scenarios appears to be sufficient to achieve a high level of accuracy.

[3] In an individual simulation the shortfall is the amount by which the fund required to provide the guarantee exceeds the value of the policyholder fund.
Shortfall on the money-back guarantee at vesting = max { £100000 – fund at maturity, 0 }
Shortfall on the GAO at vesting = max { max[£100000, fund at maturity] * (a_{OMO} / a_{GAO}) – fund at maturity, 0 }

Figure 2: Validation of Proxy Function (3 risk variables)

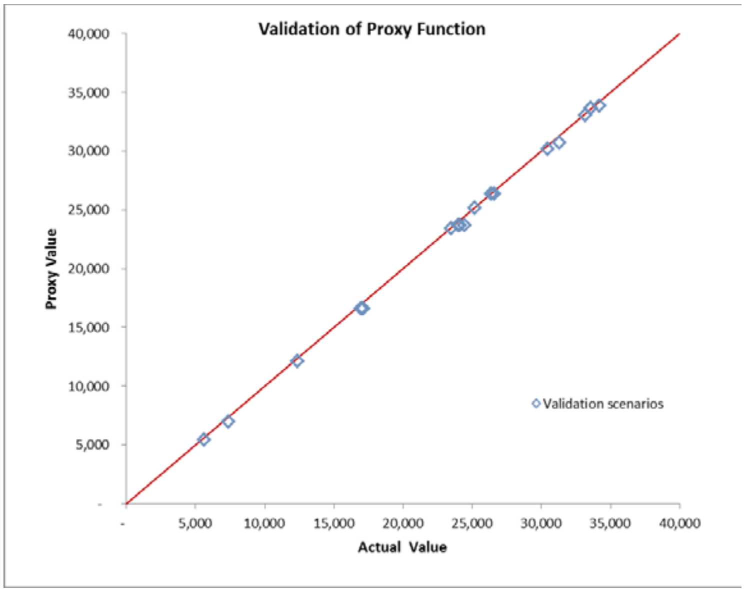

In line with expectations, the scenarios where the liability value is highest involve a significant fall in the level of the nominal yield curve. In the remainder of this paper we consider some of the practical challenges which arise when implementing proxy techniques for insurance liabilities. The first is how to extend the proxy fitting to cover non-market risks and the second looks in more detail at application of the proxy for interest rate stresses.

3. Implementation Topic 1: Modeling Demographic Risk

Given the large number of fitting scenarios involved to determine the proxy function, it is necessary to have an efficient means of running the fitting scenarios through the insurer's actuarial valuation model. For market risks this is generally straightforward as the actuarial valuation models are capable of reading in economic scenario files containing thousands of multivariate scenarios. However, the execution of thousands of demographic stresses is sometimes difficult to automate in an insurer's valuation model as it usually involves updating parameter(s) in the ALM system's assumption file. The challenge of modeling demographic risks does not lie with the theory (the technique of LSMC can be applied to non-market risks just as readily as for market risks), but lies in the practicalities of automating the processing of the fitting stresses.

To illustrate a practical approach to modeling demographic risks, the example has been extended to include an additional non-market risk variable. In the valuation model we have incorporated a constant lapse assumption during the accumulation phase prior to the annuity vesting. The liability value will vary with the lapse assumption and the proxy function can be modified to include the exposure to lapse risk. To mimic reality, it is assumed that the lapse assumption is input manually in to the actuarial valuation model and therefore, from a practical perspective, it is not possible to dynamically change the lapse assumption in the same way as performed for the market risk variables.

Recall that there were 20,000 fitting scenarios. We split the fitting scenarios in to 5 tranches. Each tranche has a different value for the lapse assumption and the initial equity allocation. Five discrete values are considered for the lapse assumption (0%, 2.5%, 5%, 7.5%, 10% per annum). Each tranche therefore contains 4,000 fitting scenarios and, within a tranche, the lapse assumption does not vary but the SOBOL sequence is still used to uniformly span the market risks. For each fitting scenario two inner simulations are produced exactly as before.

The purpose of adopting the "batching" approach is to aid with the running of the fitting scenarios and ensure that the process is practically feasible for the insurer to execute, while still incorporating non-market risks as a risk driver in the most efficient way. In this example the insurer would have to set-up five model runs, each with a different set of assumptions. The ease with which this can be achieved will depend on the insurer's individual ALM set-up.

Once the valuation results have been obtained for the fitting stresses, the B&H Proxy Generator performs the proxy fitting across the four variables. The validation results are shown below. The validation scenarios have been extended to include stresses to the demographic risks as well as market risks, and include multivariate stresses across all four risk factors.

Figure 3: Validation of Proxy Function (4 risk variables)

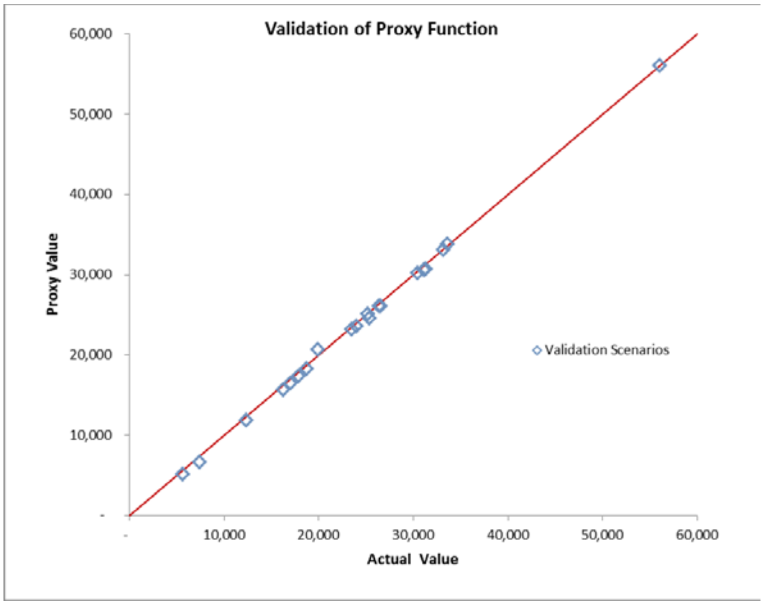

The proxy function still provides a very good fit— although there is a deterioration in the quality of the fit relative to the original example. This is not surprising given that we have maintained the same scenario budget (20,000 fitting scenarios) but increased the number of risk variables. The constraints on the maximum number of terms and the maximum order of the terms in the polynomial has been relaxed. The quality of fit can of course be improved by increasing the number of fitting scenarios.

We have essentially adopted a blended approach—LSMC for the market risks and curve fitting for the demographic risks. This example demonstrates how demographic risks can be incorporated in to proxy functions, while still adopting the LSMC approach to efficiently and accurately capture market risk. We believe this represents a sensible compromise between the superior efficiency of LSMC relative to curve fitting whilst acknowledging the practical constraints in running thousands of scenarios with different assumptions through insurer's actuarial models.

There are numerous uses for the proxy function—the primary application is often to use the proxy alongside a risk scenario generator to produce a distribution for economic capital, which may then be used to determine the capital requirement. However, the proxy is also a valuable tool for asset-liability management as it provides a means of instantly re-valuing the liabilities under stress and scenario tests. For example, when applied to a whole block of business, the business can be instantaneously re-valued to reflect a stress to the initial yield curve. However, this also presents its own challenges as discussed in the next section.

4. Implementation Topic 2: Yield Curve Modeling

This section focuses on using proxy functions in the context of interest rate modelling and re-valuing the liabilities under a specified yield curve scenario. Here the challenge is how to model a full term structure using the proxy, when the proxy typically only contains two or three risk variables representing interest rate risk. For example, if the yield curve sensitivity is expressed as a full term structure stress (i.e. multiple points on the curve), is it possible to accurately fit all points on the yield curve given that there are only two interest rate variables in the proxy function?

We commence by considering a two-factor model and then move on to the case of a three-factor model.

2-Factor Interest Rate Model

For this example we will consider the 2-factor Black-Karasinski model (2FBK) which is a 2-factor short rate model defined by:

$$d\ln\big(r(t)\big) = \alpha_1\big[\ln\big(m(t)\big) - \ln\big(r(t)\big)\big]dt + \sigma_1 dZ_1$$
$$d\ln\big(m(t)\big) = \alpha_2\big[\mu(t) - \ln\big(m(t)\big)\big]dt + \sigma_2 dZ_2$$

There are two stochastic shocks, Z_1 and Z_2. The model is calibrated by specifying the initial parameter set $\{\mu(0), \sigma_1, \sigma_2, \alpha_1, \alpha_2\}$. The evolution of the yield curve over the next time step is uniquely defined by $\{Z_1, Z_2\}$, along with the initial parameter set.

In this example we assume that the *outer fitting* scenarios used to calibrate the proxy function have been produced by using the 2FBK model to span the space of possible yield curve movements over the next year. The initial yield curve for each scenario is a realisation of the 2FBK model for a specific pair (Z_1, Z_2). By running thousands of pairs (Z_1, Z_2) through the model we simulate a set of yield curve stresses and, crucially, each yield curve is uniquely defined by a specific pair (Z_1, Z_2). This forms the set of fitting scenarios for determining the proxy function. In the proxy function we only need two variables to represent interest rate risk. These two variables may be (Z_1, Z_2), or two points on the yield curve, or two variables such as the short rate and the slope of the curve.

Proxy functions are typically used to produce a distribution for economic capital by running thousands of real-world scenarios, produced by the risk scenario generator, through the proxy function. The approach outlined above works well where the model used to produce the fitting scenarios is the same as the model for the (real-world) risk scenario generator. In our example this is the 2FBK model for both the production of fitting scenarios and the simulation of the real-world scenarios in our economic capital model. For all yield curves produced by the risk scenario generator it is possible, by construction, to directly map the scenarios produced on to the chosen two variables. In other words, for any yield curve chosen from the "universe" of yield curves produced by the 2FBK model, we can accurately value our liabilities using the proxy function. But what happens if we want to specify an arbitrary yield curve change – can we map the yield curve stress on to our two risk variables and value using the proxy? For example, what happens if we simply want to evaluate the liability proxy function for a parallel yield curve stress?

We have a proxy function which can be written as:

Liability value = $f\,(x_1, x_2)$ where x_1 and x_2 represent two yield curve factors (e.g. short and long rate)

The challenge is finding a pair (x_1, x_2) which produce a parallel yield curve stress. The analysis below looks at the example of ±100bp yield curve stresses and assesses if an exact fit can be achieved.

Using an optimization routine (such as Levenberg-Marquardt) the pair (x_1, x_2) can be determined which gives rise to the yield curve closest to the ±100bp target yield curve. The exact approach adopted to do this will depend on whether analytical formulae exist for bond prices or whether numerical methods have to be used. The optimization can therefore either focus on two individual points or minimize the sum of differences across the whole yield curve (potentially applying a weighting scheme to target the most relevant parts of the yield curve). The values for the optimal pair (x_1, x_2) can then be input to the liability proxy function to estimate the liability value under the +100bp yield curve stress.

As an example, we can consider the EUR real-world calibration at 31/12/12 of the 2FBK model and use a proxy function which has two interest rate risk variables, consistent with the 2-factor interest rate model. The graphs below show the best fit i.e. the curve produced by the pair (x_1, x_2), where x_1 and x_2 give rise to the closest fit to the target yield curve.

Figure 4: Parallel Stresses under a 2-factor structural model

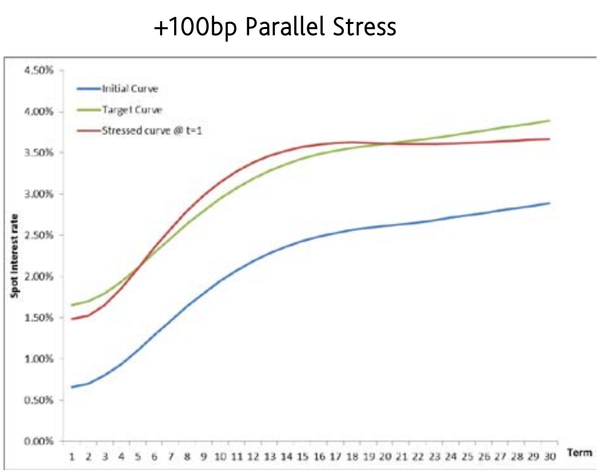

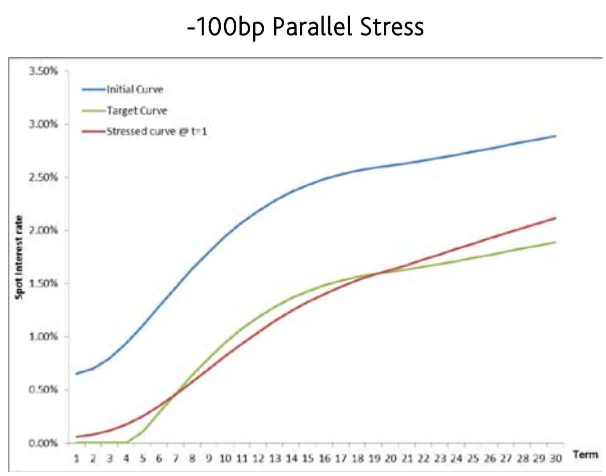

There are only two factors with which to fit a full term structure. It is not possible to get an exact fit to the target parallel yield curve stress. The ability to achieve a parallel shift is influenced by the shape of the initial yield curve (which was steeply upward sloping at 31/12/12). The calibration of the 2FBK model at 31/12/12 determines the universe of yield curves which are spanned by the model. At other dates the results would be different.

The above analysis demonstrates that, for the case where a structural 2-factor model has been adopted, parallel yield curve shifts can be very approximately replicated by using an optimization routine to determine the most appropriate values for the two risk variables in the proxy function.

3-Factor Interest Rate Model

We now consider a 3-factor model specified by:

$$StressedForwardRate(t) = InitialForwardRate(t) * \exp\{\beta_1(t)Z_1 + \beta_2(t)Z_2 + \beta_3(t)Z_3\}$$

where the β_i are arrays representing the factor loadings and Z_i are standard Brownian Motions. The model above defines a methodology for generating sensitivities to the initial yield curve, based around three components. A market-consistent interest rate model is also required to value the liabilities under each of these yield curve sensitivities.

The choice of the β_i will determine the universe of yield curves which will be spanned by the model and thus define the calibration space for the proxy function. A typical choice is to apply the technique of principal components analysis (PCA) to the historical dataset of log changes in forward rates and extract the principal components of the correlation matrix. Under this approach the first factor is typically viewed as describing shifts, the second factor as a twist factor and the third factor viewed as a curvature factor.

In this example we will use the factor loadings, β_i , as per the standard calibration of the B&H 3-factor LMM+ model. These have been derived to reproduce the real-world target correlation structure between log changes in forward rates whilst also providing an economic interpretation. The approach is to calibrate factors one and two such that they can be interpreted as short and long rate factors respectively. The third factor represents curvature.

We pose the same question as before—can we accurately map a specified yield curve stress on to the 3-factors, $\{Z_1, Z_2, Z_3\}$ and thus evaluate the proxy function under the stress. Taking the same example as before of a ±100bp stress, the graphs below show the best fit i.e. the curve produced by the pair (Z_1, Z_2, Z_3), where Z_1, Z_2 and Z_3 give rise to the closest fit to the target yield curve.

Figure 5: Parallel Stresses under a 3-factor structural model

+100bp Parallel Stress　　　　-100bp Parallel Stress

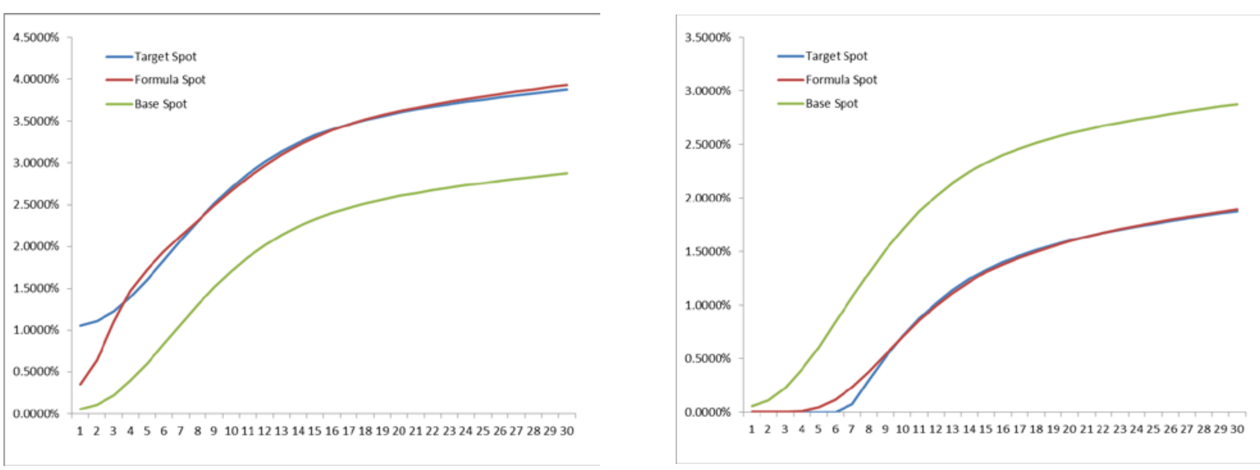

In this case the fit to the target parallel yield curve stresses is very good. Using a three factor model to calibrate the proxy function enables a wider range of yield curves to be fitted. In this simplistic case we considered a parallel shift – indeed, if the first factor loading did represent a perfect parallel shift, then this would ensure that we could exactly replicate a parallel shift. However, caution must be exercised when deliberately constructing the factor loadings, β_i, to target specific yield curve stresses. It is desirable to have economically intuitive factors which combine to produce a set of economically viable yield curves.

Figure 6 below considers some more complex examples such as yield curve twists (here around the 10-year point on the curve):

Figure 6: Yield Curve twists under a 3-factor structural model

Short End up, Long End Down

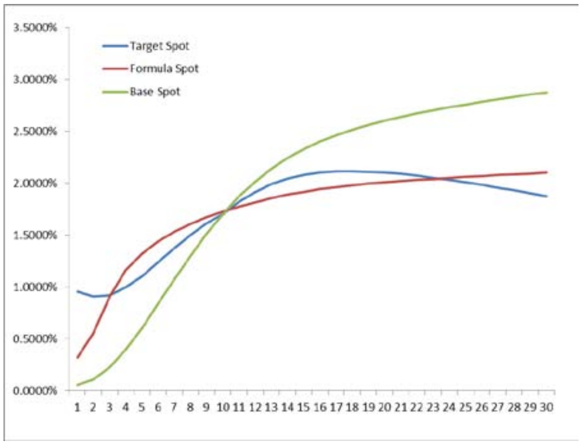

Short End Down, Long End Up

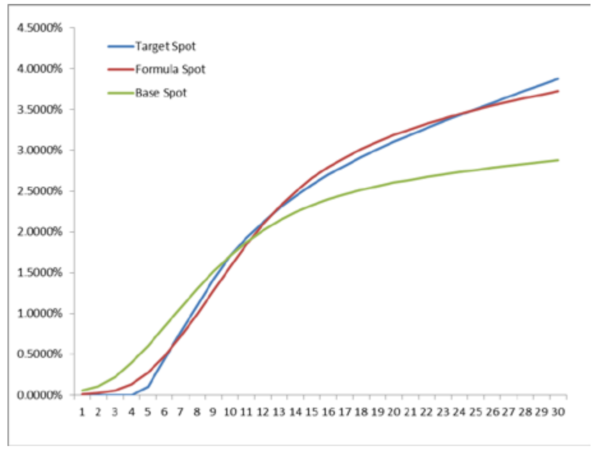

Three factors provide extra flexibility to consider a wider range of yield curve shapes but, as shown in the graphs above, some yield curve shapes are still challenging—this is a direct consequence of the dimension reduction in the modeling.

An insurer may wish to re-value liabilities under yield curve stresses that are not consistent with the yield curve model used to produce 1-year simulations of yield curve changes. This section has shown that this issue can be addressed by mapping yield curve stresses to two or three factors which can then be input in to a liability proxy function.

5. Conclusions

In this paper we have considered a typical insurance liability with a complex embedded option, such that the policy has exposure to both market and demographic risks. Using the technique of LSMC a proxy function has been determined for the cost of the guarantee and validation has been performed to assess the quality of fit of the proxy function. Proxy fitting involves reducing the dimensionality of the underlying risk factors. This selection of the risk drivers, along with the fitting range for the risk drivers, is critical as it influences both the quality of the proxy function and the potential uses of the proxy function. The paper considered two of the practical challenges which insurers encounter in implementing proxy functions – namely incorporating demographic risk factors and also yield curve mapping. These challenges arise for different reasons but both can be resolved using the approaches outlined in this paper.