MOODY'S ANALYTICS

B&H RESEARCH
NOVEMBER 2013

Proxy functions for the projection of Variable Annuity Greeks

Aubrey Clayton PhD
Steven Morrison PhD
Craig Turnbull FIA
Naglis Vysniauskas
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

Overview

Variable Annuity providers in North America and Europe must perform stochastic projections of the behaviour of their dynamic hedging programs over the lifetime of these long-term liabilities in order to obtain recognition for the risk mitigation benefits of the hedging in their regulatory capital assessments. The computational demands of this calculation have proven to be a significant hurdle to firms. As a result, many firms are not able to obtain realistic levels of capital relief or are undertaking enormous complex nested stochastic calculations that are expensive, unwieldy and that may involve arbitrary simplifications that undermine confidence in their results. We believe this paper breaks new ground by introducing an entirely different methodology for implementing the highly demanding modelling required in this area; an approach that is significantly more efficient, accurate and objective than those applied in industry up until now.

Contents

1. Introduction

Variable annuity writers in North America and Europe must perform real-world stochastic projections over the lifetime of these long-term liabilities in their regulatory capital assessments in order to obtain recognition for the risk mitigation benefits of any hedging program in this assessment. Variable annuity Greeks need to be estimated in each real-world scenario, at each future point in time where the hedge portfolio is rebalanced.

The standard approach to estimation of such Greeks involves estimating market-consistent values before and after the corresponding risk factor is 'bumped', with each valuation being carried out using risk-neutral Monte Carlo simulation. This creates a nested simulation problem, with a large number of 'inner' risk-neutral simulation scenarios emerging from each time-step in each 'outer' real-world scenario. The total number of inner simulations required to accurately model a real hedge program using this approach is likely to be unfeasibly large for most firms given current model run-time. Any practical nested simulation approach is likely to adopt simplifying assumptions in order to reduce run-time to acceptable levels, for example a significant reduction in the frequency at which the hedge portfolio is assumed to rebalance.

In recent years, the use of 'proxy functions' has emerged as a practical alternative to nested simulation, and is now widely used by insurance firms in the projection of market-consistent value of liabilities for the purpose of calculation of 1-year VaR. These functions provide analytical approximations to market-consistent value that can be evaluated quickly compared to a full risk-neutral simulation. While previous use of proxy functions has tended to focus on projection of market-consistent value at a single future time horizon, recent research has demonstrated how the technique extends to projection of market-consistent value across entire paths over the liability lifetime (Morrison, Turnbull and Vysniauskas 2013). In this note, we take this idea one step further and develop proxy functions for various Greeks in addition to market-consistent value. The availability of such proxy functions provides a practical alternative to full nested simulation, allowing dynamic hedge programs to be projected at realistic (e.g. daily) frequency.

2. Example liability and valuation model

In this note we consider the specific problem of projection of the market-consistent value, Delta, Rho and Vega for a 10-year forward at-the-money vanilla put option on an equity index.[1] To estimate the market-consistent value of this liability, we assume a risk-neutral valuation model in which equity excess log-returns are normal with constant volatility, while risk-free interest rates follow the Barrie & Hibbert LMMPlus model. While a constant volatility model may appear unrealistically simple, such models are often used to value VA guarantees in practice. Furthermore, for the particular liability considered here, market-consistent value and Greeks can be calculated using standard Black-Scholes analytical formulae.[2] In the current exercise, such analytical values are useful as they provide a benchmark against which we can validate fitted proxy functions.

3. Fitting methodology for multi-period Greeks

All proxy fits described in this note are based on the Least Squares Monte Carlo (LSMC) approach, whereby proxy functions are fitted to Monte Carlo estimates using least squares to perform the fit (Bauer, Bergmann and Reuss 2010). We have investigated three distinct approaches of using the LSMC technique to fit multi-period Greeks, as described below.

Method 1: Estimate values → Fit value proxy function → Differentiate

Our first method is a relatively straightforward generalization of the multi-step proxy functions previously developed for market-consistent value in (Morrison, Turnbull and Vysniauskas 2013). As described in that paper, we fit a set of polynomial proxy functions $V_t^{proxy} = V_t^{proxy}(S, R, \sigma)$ using least squares regression i.e. given a set of Monte Carlo estimates of market-consistent value $\hat{V}_{i,t}$ (corresponding to fitting scenario i at time t, with equity index $S_{i,t}$, risk-free spot rate $R_{i,t}$ and volatility $\sigma_{i,t}$) we choose the polynomial coefficients to minimize the sum of squared residuals over all fitting scenarios:

$$\sum_i \left(\hat{V}_{i,t} - V_t^{proxy}(S_{i,t}, R_{i,t}\sigma_{i,t})\right)^2$$

We repeat this fitting process for each time t of interest.

Now the Greeks are defined as sensitivities of value to changes in the various risk factors. Thus we can then simply *differentiate* the value proxy functions to estimate proxy functions for the Greeks:

$$\Delta_t^{proxy}(S, R, \sigma) = \frac{\partial}{\partial S} V_t^{proxy}(S, R, \sigma)$$

$$\rho_t^{proxy}(S, R, \sigma) = \frac{\partial}{\partial R} V_t^{proxy}(S, R, \sigma)$$

$$\mathcal{V}_t^{proxy}(S, R, \sigma) = \frac{\partial}{\partial \sigma} V_t^{proxy}(S, R, \sigma)$$

Method 2: Estimate Greeks → Fit Greeks proxy functions (global regression)

Although the proxy functions developed in Method 1 may provide an acceptable fit to market-consistent value, we may observe some divergence between estimated and actual *sensitivities*. The polynomial function for value may include some higher order terms, and even a small error in the estimated coefficients of those terms will become magnified when taking derivatives. As a result, we have considered an alternative method, whereby we estimate the Greeks directly and fit proxy functions to these estimates.

One relatively simple method for estimating Greeks directly is using a 'bump and revalue' approach.[3] This involves bumping the relevant risk factor by a small amount δx, revaluing the option and estimating the derivative (sensitivity) as the finite difference

[1] Here we define Rho as sensitivity to the risk-free spot rate corresponding to the liability maturity, and Vega as sensitivity to the implied volatility corresponding to the liability maturity and strike.

[2] Note that the presence of stochastic interest-rates means that the Black-Scholes formulae will only be approximate here.

[3] See (Cathcart, McNeil and Morrison 2011) for a comparison of various methods for the estimation of variable annuity Greeks.

$\frac{\partial V}{\partial x} \cong \frac{V(x+\delta x)-V(x)}{\delta x}$. The difficulty with this approach is that when we bump the risk factor we have to create an additional scenario, and in the current context we potentially have to do this at a (large) number of different points in time. This has the potential to significantly increase the number of required scenarios compared to the previous method. Indeed, in this way we end up performing a nested simulation (albeit with a smaller number of inner scenarios than would be required for a 'full' nested simulation).

To avoid this problem, we have reused the same 'bump' scenarios at different time-steps. We do this by creating 'clusters' of scenarios, as follows:

- Firstly we create a set of independent 'base' scenarios. Each base scenario has different initial parameters for the key risk factors, in this case, the initial equity index level, yield curve level, and equity volatility.
- For each base scenario we bump (or stress) a single risk factor at t=0, while leaving all other risk factors unchanged. We then regenerate the scenario under this stressed initial condition (but reuse the same random numbers to create the scenario). We do this separately for each risk factor: initial equity index level, yield curve level, and equity volatility.
- A cluster is then defined as a set of four scenarios: base plus the three different stresses. All scenarios within a particular cluster can be expected to be 'close' to each other in risk factor space, at all future points in time, provided the initial stresses are small enough and the same random scenarios are used for all scenarios in the cluster. Note however that the difference in risk factors at any future point in time will not in general be the same as the assumed stress at t=0, and in particular a risk factor may be identical in the base and stressed scenarios at t=0 but may diverge in future as the simulation progresses.[4]

Given a particular cluster i at time t, we can estimate the three Greeks $\hat{\Delta}_{i,t}, \hat{\rho}_{i,t}, \hat{\mathcal{V}}_{i,t}$ by solving the following system of linear equations:

$$\begin{pmatrix} S_{i,t}^{stress\ 1} - S_{i,t}^{base} & R_{i,t}^{stress\ 1} - R_{i,t}^{base} & \sigma_{i,t}^{stress\ 1} - \sigma_{i,t}^{base} \\ S_{i,t}^{stress\ 2} - S_{i,t}^{base} & R_{i,t}^{stress\ 2} - R_{i,t}^{base} & \sigma_{i,t}^{stress\ 2} - \sigma_{i,t}^{base} \\ S_{i,t}^{stress\ 3} - S_{i,t}^{base} & R_{i,t}^{stress\ 3} - R_{i,t}^{base} & \sigma_{i,t}^{stress\ 3} - \sigma_{i,t}^{base} \end{pmatrix} \begin{pmatrix} \hat{\Delta}_{i,t} \\ \hat{\rho}_{i,t} \\ \hat{\mathcal{V}}_{i,t} \end{pmatrix} = \begin{pmatrix} \hat{V}_{i,t}^{stress\ 1} - \hat{V}_{i,t}^{base} \\ \hat{V}_{i,t}^{stress\ 2} - \hat{V}_{i,t}^{base} \\ \hat{V}_{i,t}^{stress\ 3} - \hat{V}_{i,t}^{base} \end{pmatrix}$$

Having estimated the Greeks in this way we can then perform three *separate* least squares fits to estimate polynomial proxy functions for each of the Greeks, $\Delta_t^{proxy}, \rho_t^{proxy}, \mathcal{V}_t^{proxy}$, i.e., choose the polynomial coefficients to minimize:

$$\sum_i \left(\hat{\Delta}_{i,t} - \Delta_t^{proxy}(S_{i,t}, R_{i,t}\sigma_{i,t})\right)^2$$

$$\sum_i \left(\hat{\rho}_{i,t} - \rho_t^{proxy}(S_{i,t}, R_{i,t}\sigma_{i,t})\right)^2$$

$$\sum_i \left(\hat{\mathcal{V}}_{i,t} - \mathcal{V}_t^{proxy}(S_{i,t}, R_{i,t}\sigma_{i,t})\right)^2$$

We can of course also estimate a proxy function for the market-consistent value by regressing on the base scenario values i.e. choose the coefficients of V_t^{proxy} to minimize:

$$\sum_i \left(\hat{V}_{i,t}^{base} - V_t^{proxy}(S_{i,t}, R_{i,t}\sigma_{i,t})\right)^2$$

Comparing with Method 1, we see that the same type of least squares approach is used to perform proxy fits, but that separate fits are carried out to value, Delta, Rho and Vega estimates, rather than performing a single fit to value and differentiating to calculate the Greeks.

[4] For example, a stress to the initial yield curve will result in different future equity scenarios even if the initial equity index is unstressed, since equity returns depend on risk-free interest rates.

Method 3: Estimate Greeks $\rightarrow$ Fit Greeks proxy functions (local regression)

The above methods attempt to approximate liability values and Greeks using polynomials over the entire risk factor space. This is not always appropriate for certain shapes of function. For example, as the time-to-maturity of a put option tends to zero the shape of the Delta, as a function of the underlying equity price, approaches a step function, which is not well described by a low order polynomial.

The difficulty in fitting such shapes motivates the use of alternative forms of proxy function which are less restrictive in their shape. In this note we have used a type of 'local regression', whereby we fit local polynomial functions at each point in the risk factor space, using a weighted least squares technique with more weight on placed on 'nearby' fitting points (Cleveland and Devlin 1988). For example, if we want to estimate the option value at time t, where the risk factors take the values (S, R, σ), we evaluate using a polynomial V_t^{proxy} whose coefficients have been chosen so as to minimize the *weighted* sum of squared residuals:

$$\sum_i w\left(\frac{S_{i,t}-S}{h^S}, \frac{R_{i,t}-R}{h^R}, \frac{\sigma_{i,t}-\sigma}{h^\sigma}\right)\left(\hat{V}_{i,t}^{base} - V_t^{proxy}(S_{i,t}, R_{i,t}, \sigma_{i,t})\right)^2$$

Note that the weight $w\left(\frac{S_{i,t}-S}{h^S}, \frac{R_{i,t}-R}{h^R}, \frac{\sigma_{i,t}-\sigma}{h^\sigma}\right)$ depends on the distance of the fitting points from the point at which we wish to evaluate the proxy function, and the weight function is typically chosen so as to put less weight on fitting points that are further away from the evaluation point.[5] Since the fitting space is adjusted dynamically to reflect a small neighborhood around the chosen point, we can accurately estimate the function in that neighborhood using only low-order polynomials, for example linear functions or even constants.

This local regression technique is a relatively straightforward extension of the 'global' regression adopted in Methods 1 and 2. Indeed the global regression can be considered a special case with equal weighting on all fitting points, and thus the proxy function is a single polynomial applied globally. However, more generally we can describe the proxy function using a different polynomial at every point in risk factor space. This clearly gives us much more flexibility in shape (and hence potential for increased fitting quality) with the potential drawback that the proxy function is more complex, harder to communicate and computationally more expensive.

Interpolation in time

So far, we have described how to fit proxy functions at a particular point in time (for example, V_t^{proxy} denotes the value proxy function at time t). Effectively, we fit a different set of proxy functions for each time of interest. However, in the context of hedge projection we need proxy functions at each time at which the hedge portfolio is rebalanced, and in practice this could be as frequently as daily. In such cases, it would be computationally costly to proxy functions at every required time. Rather, it would be convenient to perform estimates at a much smaller frequency (e.g. annual) and use an interpolation scheme to fit proxy functions at intermediate times.

One way of interpolating in time is simply to include time as a variable in the proxy function, $V^{proxy} = V^{proxy}(t, S, R, \sigma)$. For example, in the case of the general local regression (Method 3), if we want to estimate the option value at time s, where the risk factors take the values (S, R, σ), we choose the local polynomial fit by minimizing the following sum of squared residuals:

$$\sum_{i,t} w\left(\frac{t-s}{h^S}, \frac{S_{i,t}-S}{h^S}, \frac{R_{i,t}-R}{h^R}, \frac{\sigma_{i,t}-\sigma}{h^\sigma}\right)\left(\hat{V}_{i,t}^{base} - V^{proxy}(t, S_{i,t}, R_{i,t}, \sigma_{i,t})\right)^2$$

4. Case study: Fits and Validations

Having fitted proxy functions using the three methods described above, we evaluate quality of fit by calculating Greeks in a number of out-of-sample validation scenarios, and comparing proxies with actual Greeks calculated using analytical (Black-Scholes) formulae.[6] Validation scenarios were chosen according to an equally spaced grid spanning the range of 1st to 99th percentile of the real world distribution for each risk factor. In this way, we are sure to validate against an extreme range of all

[5] h^S, h^R, h^σ are known as bandwidth parameters and control how quickly the weights decay as a function of the distance from the evaluation point.

[6] Recall that the Black-Scholes formulae valuations are only approximations to the actual value due to the presence of stochastic interest-rates.

Greeks. At time 9, the mesh of this grid was made finer in the equity index dimension, to show a more complete range of option moneyness values.

Note that the three different fitting methods described above make use of Monte Carlo scenarios in different ways, so to evaluate the methods on an equal basis we assume a fixed budget of fitting scenarios, measured by the number of valuations of the option liability at maturity. The fitting points for each method were constructed by the technique described in (Morrison, Turnbull and Vysniauskas 2013): identifying a desired range of each risk factor at time 0 (initial equity index, equity volatility, and two principal components of the initial yield curve), applying a low discrepancy sequence to efficiently span this 4-dimensional space, and then generating a single risk-neutral path to serve both as a valuation scenario and fitting point for later time-steps.

Figures 1, 2 and 3 compare proxy vs. actual Deltas for all three methods at times 1, 5 and 9 years, respectively.

Figure 1: Delta at year 1: proxy vs. actual in 216 validation scenarios

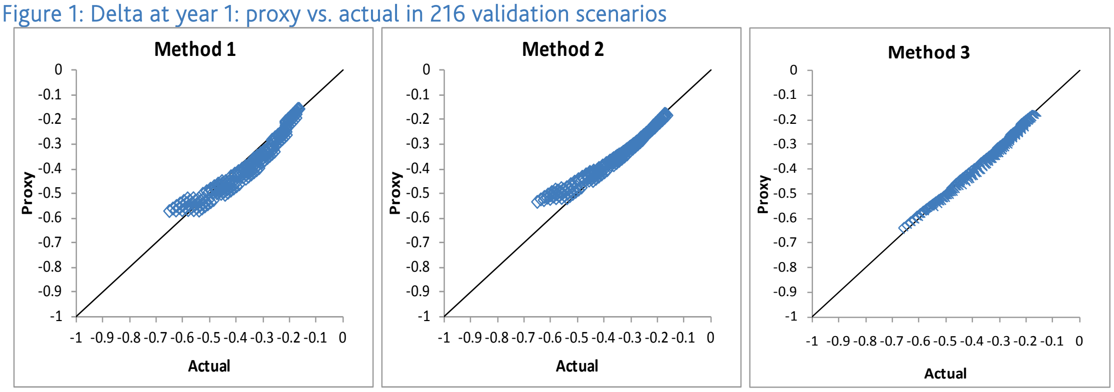

Figure 2: Delta at year 5: proxy vs. actual in 216 validation scenarios

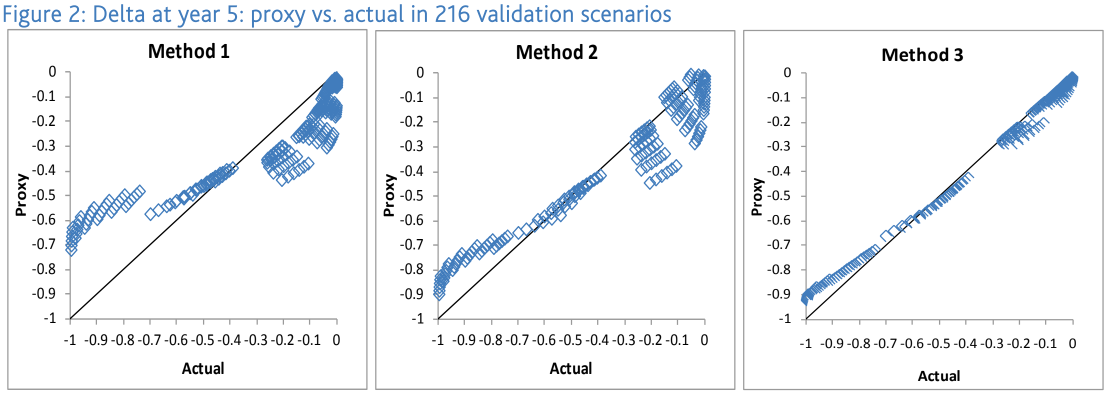

Figure 3: Delta at year 9: proxy vs. actual in 756 validation scenarios

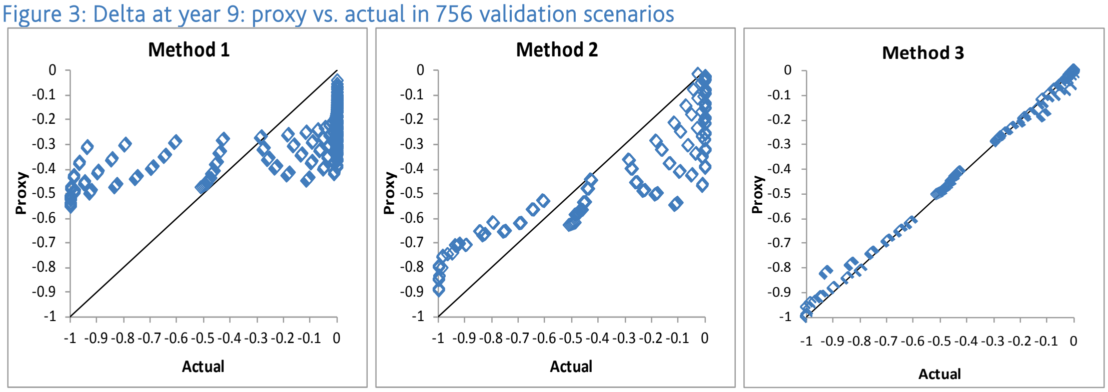

Figures 4, 5 and 6 compare proxy vs. actual Rhos for all three methods, at times 1, 5 and 9 years, respectively.

Figure 4: Rho at year 1: proxy vs. actual in 216 validation scenarios

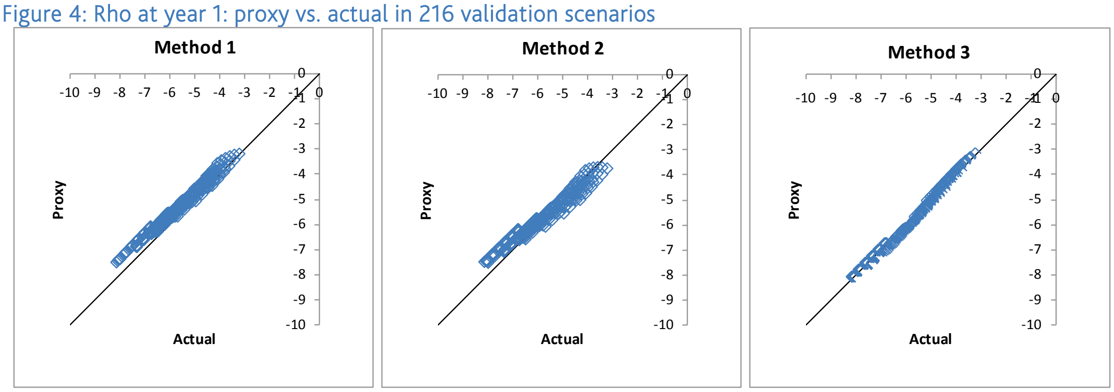

Figure 5: Rho at year 5: proxy vs. actual in 216 validation scenarios

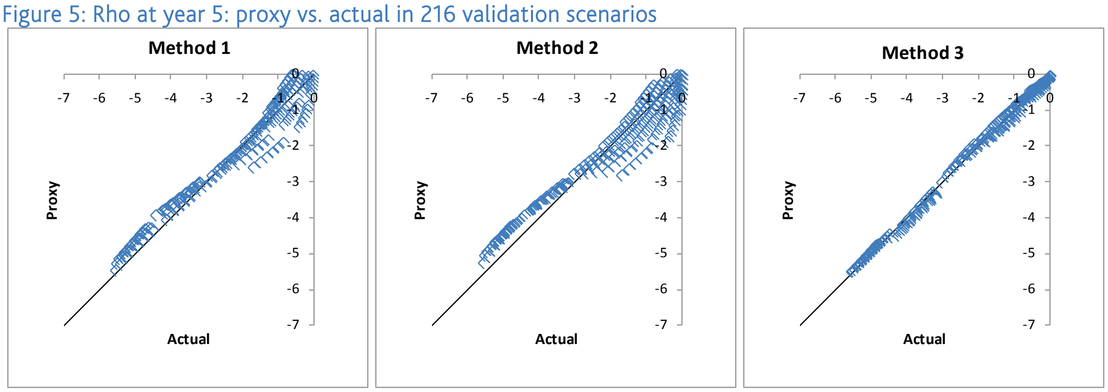

Figure 6: Rho at year 9: proxy vs. actual in 756 validation scenarios

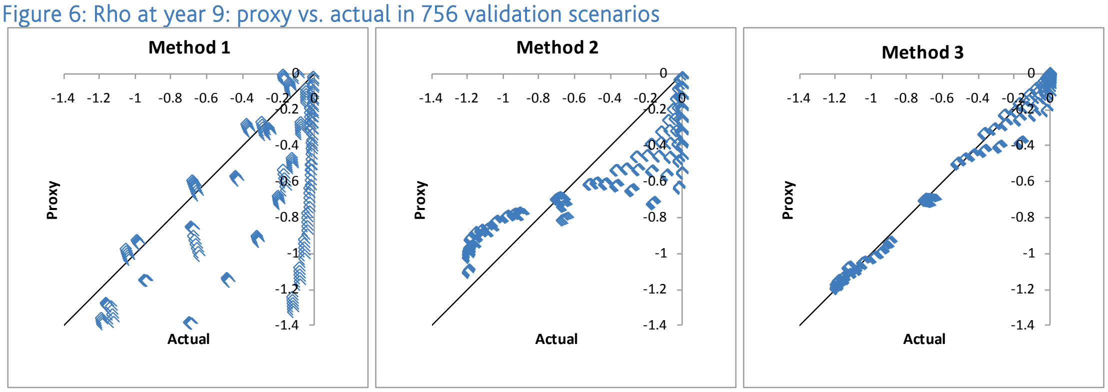

Finally, figures 7, 8 and 9 compare proxy vs. actual Vegas for all three methods, at times 1, 5 and 9 years, respectively.

Figure 7: Vega at year 1: proxy vs. actual in 216 validation scenarios

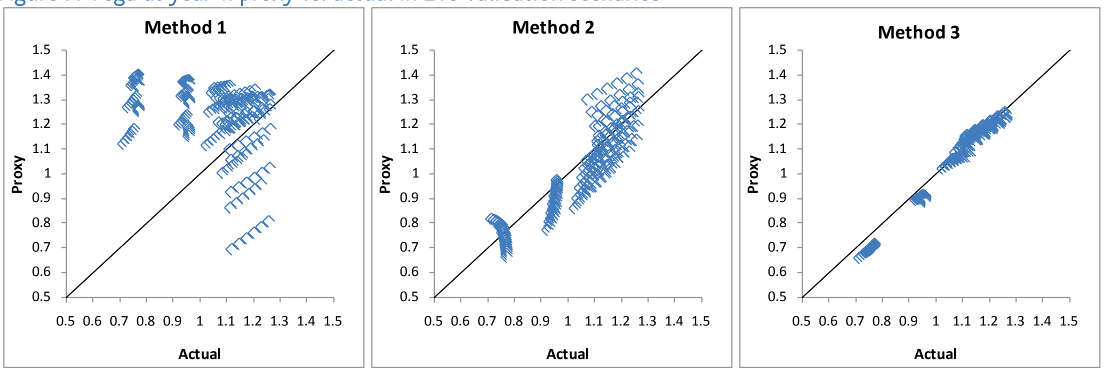

Figure 8: Vega at year 5: proxy vs. actual in 216 validation scenarios

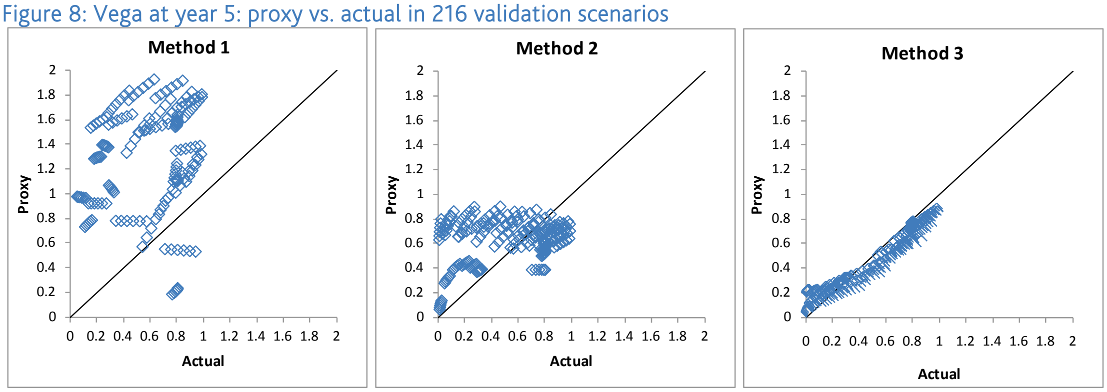

Figure 9: Vega at year 9: proxy vs. actual in 756 validation scenarios

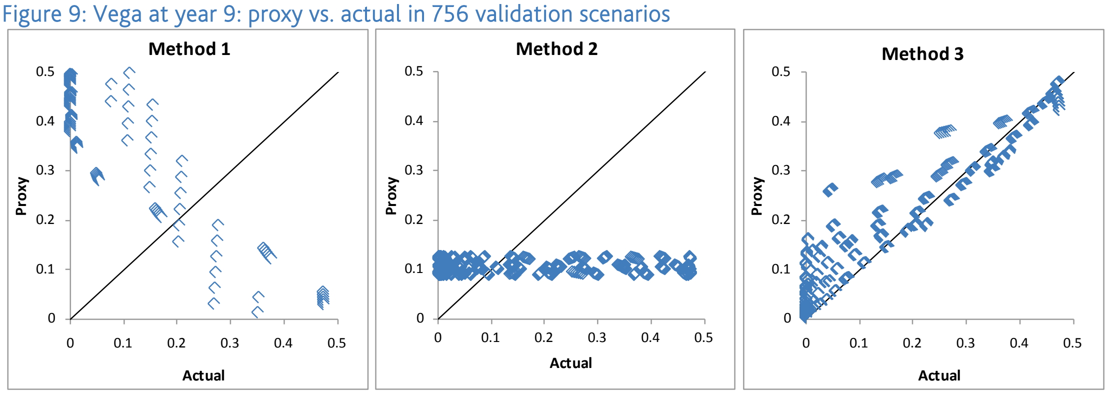

Figures 1-6 (leftmost charts) show that the Method 1 proxy functions for Delta and Rho validate well at year 1, but the quality of fit deteriorates significantly as we project forward and the liability approaches maturity. Method 1 proxy functions for Vega (Figures 7-9) are very poor at all times. The poor quality of fit observed for all Greeks at later times suggest that the Method 1 proxy functions produced here are of insufficient quality to accurately project a hedge portfolio.

Method 2 (middle charts in Figures 1-9) was introduced in an attempt to produce better Greek fits by fitting to the Greeks directly. However, in this example, Method 2 appears to offer only marginal improvement in quality of fit to Greeks, and potentially a slight deterioration in quality of fit to value (as the number of fitting scenarios dedicated to fitting value is one quarter of that used in Method 1).

Method 3 (rightmost charts in Figures 1-9) *show a vastly improved quality of fit* when the regression methodology is changed from global to local regression, particularly as the liability approaches maturity. This is intuitive as the shapes of these functions are not well represented *globally* by low order polynomials. In contrast, the local regression methodology provides a far more flexible description of the proxy function that is able, in general, to accurately fit these shapes.

As discussed in Section 3, in order to project a hedge portfolio we need to be able to evaluate Greeks at every time where the portfolio is rebalanced, and an interpolation methodology was introduced based on including time as a variable in the proxy function. To test this interpolation method, Figure 10 compares proxy and actual values and Greeks at time 4.5, using Method 3. This, and similar exercises at other interpolated times, indicate that a similar quality of fit is observed at interpolated times as at integer times (where we have explicit fitting points).

Figure 10: Proxy vs. actual values and Greeks, in 216 validation scenarios, at year 4.5 (Method 3)

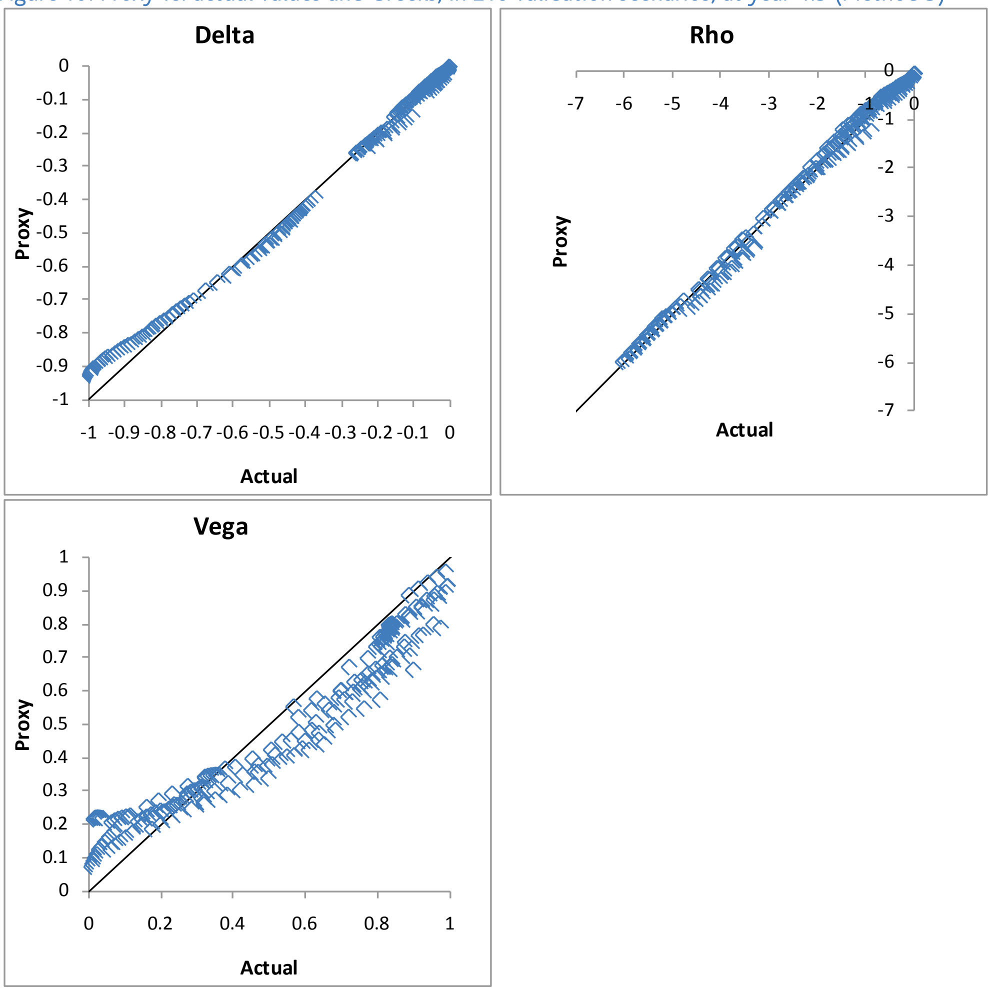

The above analysis highlights that multi-period Greeks fitting is a demanding statistical / computational challenge. To adequately describe the behavior of Greeks at various projection horizons, it is likely that more sophisticated fitting techniques than the 'vanilla' LSMC method used for 1-year VaR will be required. The above discussion suggests that local regression methods can provide the flexibility required to meet this task in a much more efficient way than through a full-blown nested stochastic implementation.

5. Hedge portfolio projection

The previous section focused on validating the quality of the Greek estimates at a number of specific projection horizons. This section considers the behavior of the *paths* produced for the Greeks over the lifetime of the product, and the impact these have on hedge portfolio behaviour. We focus specifically on the fits produced by Method 3, which we concluded in the previous section was the most promising method for multi-period Greeks fitting. We consider two illustrative 10-year daily-timestep real-world simulations[7] – one where the option expires out-the-money, and one where it matures in-the-money.

In our illustrative hedge portfolio simulation, the portfolio is assumed to consist of cash, equity index futures, a 1-year at-the-money equity option and a zero-coupon bond with maturity equal to the maturity of the option. The amounts invested in each of these assets are rebalanced on a daily basis to best match the Delta, Rho and Vega Greek estimates.

Illustrative Simulation 1: Option matures out-the-money

In Illustrative Simulation 1, the equity index performs well over the 10-year projection horizon and finishes above the forward at-the-money option strike of 1.18, and the put option therefore expires worthless. Figure 11 shows the daily time-step path for the equity index, market-consistent liability value and the dynamic hedge portfolio throughout the 10 year lifetime of the option. Two projections of the hedge portfolio are plotted: the one produced using the exact Greeks (recall that these are analytically available for the option in question), and one that is rebalanced according to the Greeks produced by the LSMC proxy fits.

Figure 11: Illustrative Simulation 1: Equity Index and Hedge Portfolio Projection

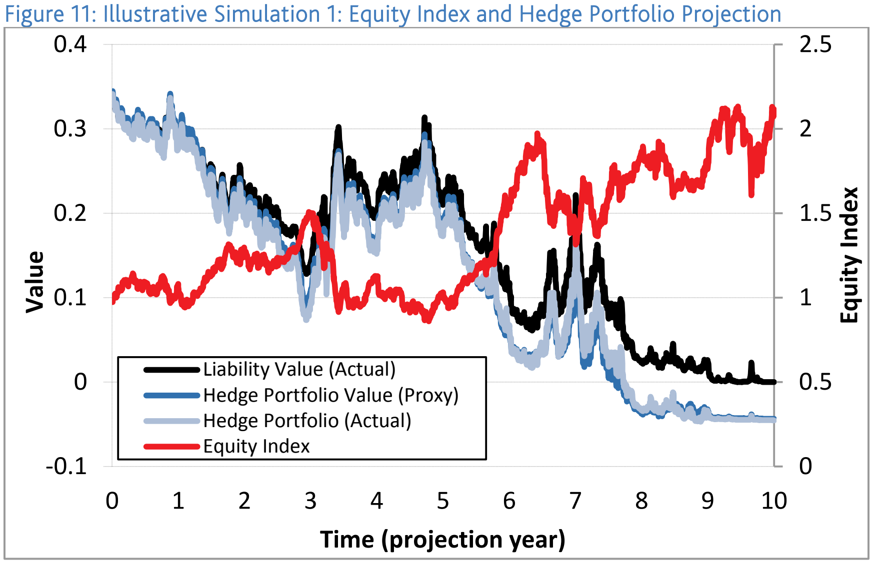

You can see from Figure 11 that the hedge portfolios do a reasonable job of matching the market-consistent liability value, though they have accumulated some hedging error by the maturity of the option. However, the comparison of the two hedge portfolios suggests that the accumulated hedging error is primarily due to the limitations of the hedging strategy rather than due to the impact of Greek mis-estimation. Given we are using a stochastic volatility with jumps equity model to simulate these paths, the hedging limitations will include jump risk and some vega mismatch risk (i.e. using the 1-year option to hedge 10-year option Vega.

The following three charts compare the paths for Delta, Rho and Vega respectively that have been produced analytically and with the proxy fits.

[7] The B&H ESG's Stochastic Volatility Jump Diffusion Model is used to produce these real-world simulations.

Figure 12: Illustrative Simulation 1: Delta Projection

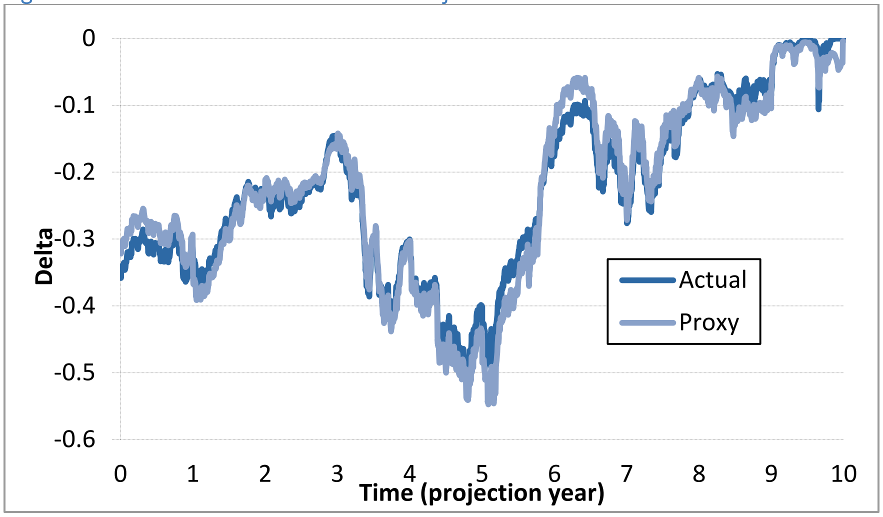

Figure 13: Illustrative Simulation 1: Rho Projection

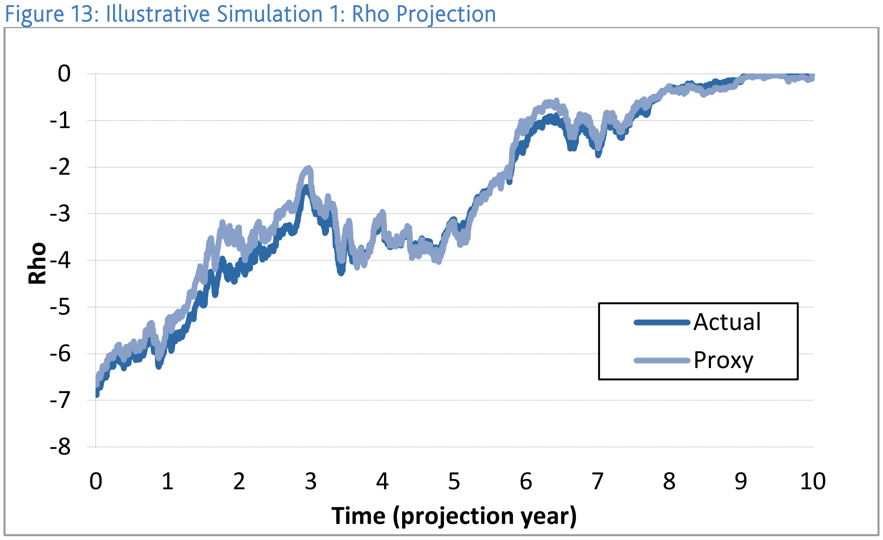

Figure 14: Illustrative Simulation 1: Vega Projection

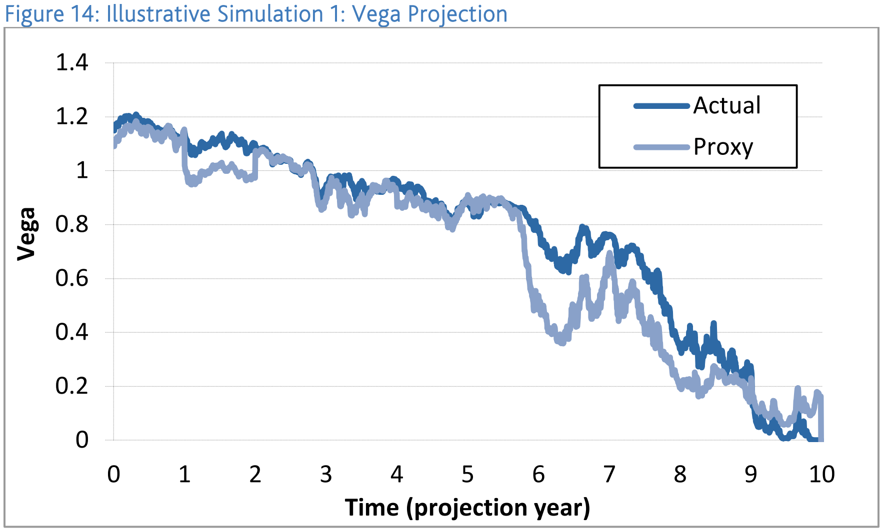

The above results are consistent with section 4 – the Delta and Rho fits are very accurate throughout the projection horizon; there is more material estimation error in the Vega fits, though they again capture the directional changes reasonably well.

Illustrative Simulation 2: Option matures in-the-money

We now move on to Illustrative Simulation 2. In this simulation, the equity index performs poorly over the 10-year horizon, and the option finishes substantially in-the-money. Figure 15 shows the simulated equity index, market-consistent liability value and hedge portfolio values produced by this simulation. In this case, you can see that the hedge portfolio has accurately tracked the liability value throughout the projection horizon, and this is the case with the proxy Greeks as well as the actual Greeks.

Figure 15: Illustrative Simulation 2: Equity Index and Hedge Portfolio Projection

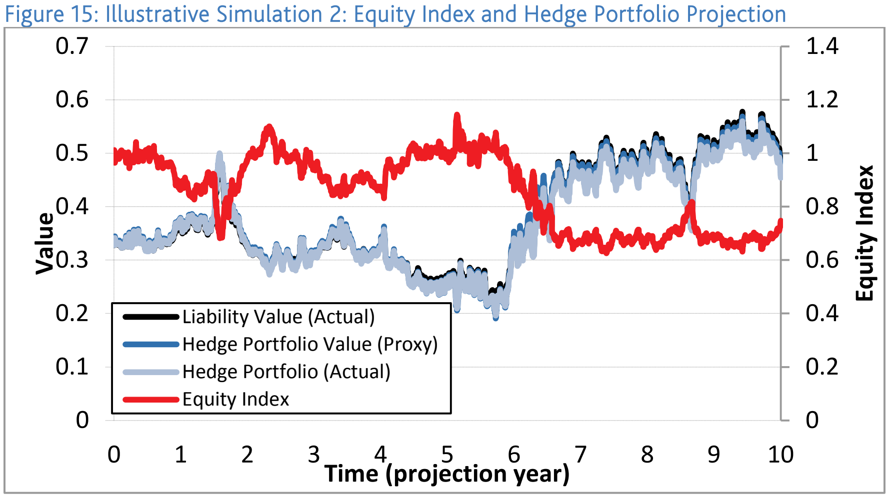

The following three charts show the projected paths for Delta, Rho and Vega respectively for Illustrative Simulation 2.

Figure 16: Illustrative Simulation 2: Delta Projection

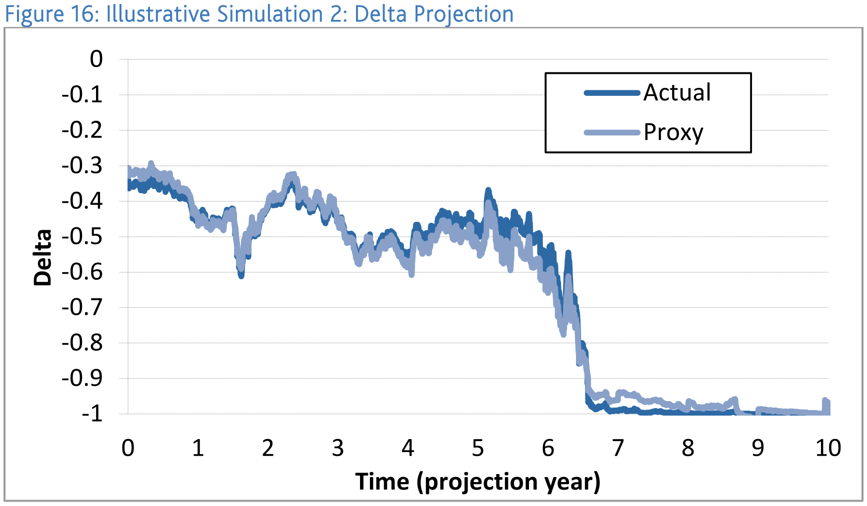

Figure 17: Illustrative Simulation 2: Rho Projection

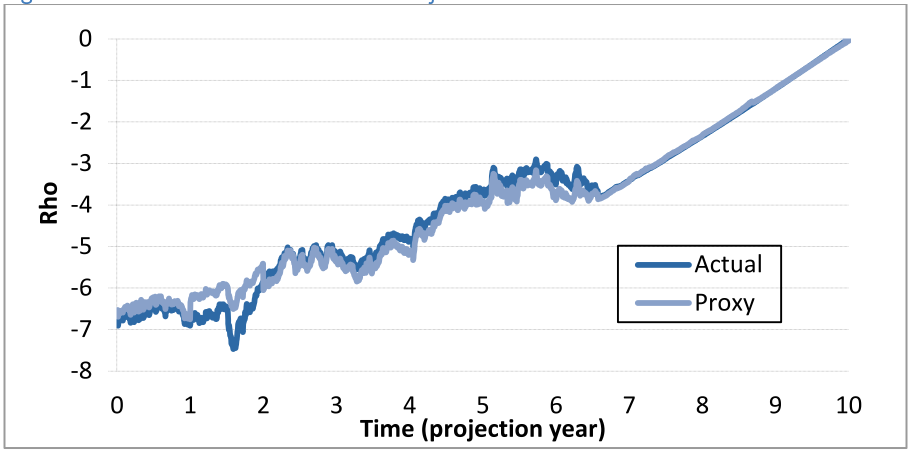

Figure 18: Illustrative Simulation 2: Vega Projection

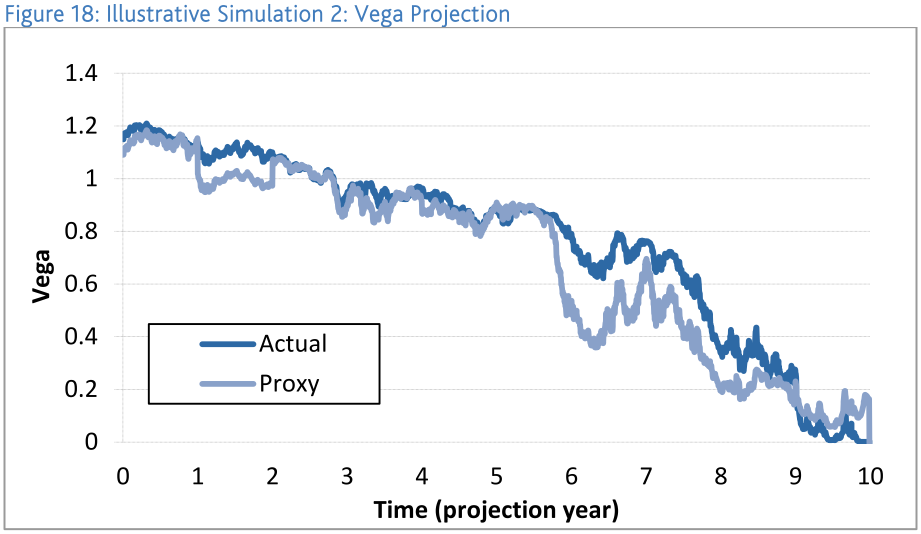

These results are again consistent with section 4: Delta and Rho behavior is very accurately fitted throughout the projection path, whilst Vega has a more material fitting error, whilst still capturing the general movements well.

6. Summary and conclusions

This paper has discussed one of the most demanding statistical computation challenges in market risk modeling for the global life industry today: stochastic projection of the dynamic hedge portfolios backing complex long-term liability guarantees such as those found in Variable Annuity business.

From the research and analysis presented in this paper, we believe that the current industry practices of performing full nested stochastic calculations with heavily 'model-pointed' liability representation can be substantially improved upon through the use of Least Squares Monte Carlo (LSMC), both in terms of statistical efficiency and accuracy and also in terms of objectivity of implementation.

However, the complexity of this modeling task means that a more sophisticated LSMC implementation methodology is likely to be required to simultaneously produce reliable estimates for a range of Greeks throughout a long-horizon path. In particular, the research in this paper has found that using local or non-parametric regression methods can provide the required flexibility and power to meet these demands and improve on the current nested stochastic implementation approaches.

References

Bauer, D, D Bergmann, and A Reuss. *Solvency II and Nested Simulations - a Least-Squares Monte Carlo Approach.* Univeristy of Ulm working paper, 2010.

Cathcart, M J, A J McNeil, and S Morrison. "Calculating Variable Annuity Liability 'Greeks' Using Monte Carlo Simulation." 2011. http://arxiv.org/abs/1110.4516#!

Cleveland, W S, and S J Devlin. "Locally-Weighted Regression: An Approach to Regression Analysis by Local Fitting." *Journal of the American Statistical Association*, 1988: 83 (403): 596-610.

Morrison, Steven, Craig Turnbull, and Naglis Vysniauskas. *Multi-year projection of market-consistent liability valuations.* 2013.