MOODY'S ANALYTICS

RESEARCH
JANUARY 2015
INSURANCE ERS

Steven Morrison PhD
Laura Tadrowski PhD
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

Quantifying the effect of dynamic hedging on 1-year VaR capital

Overview

In this note, we consider some of the technical challenges and solutions in adapting internal models to account for the effect of dynamic hedging strategies in reducing 1-year VaR capital requirements.

The 'single period' models commonly adopted to assess 1-year VaR are designed to capture risk at a single (one year) time horizon and therefore do not naturally extend to include the effects of dynamic asset (or liability) management strategies over the year. Capturing the effect of dynamic hedging raises new technical challenges to the builders of internal models, in particular:

1. Generation of economic scenarios for realistic daily risk factor paths.
2. Fast evaluation of liability Greeks.

We illustrate how these challenges can be tackled using a case study, based on a variable annuity business with a daily delta hedging strategy. Liability Greeks are approximated using accurate proxy function techniques, calibrated using Least Squares Monte Carlo (LSMC).

This case study highlights the sensitivity of hedge performance to the realised daily volatility of underlying risk factors, and the resulting sensitivity of 1-year VaR capital requirements to the assumed variation in volatility. A relatively simple Geometric Brownian Motion (GBM) model severely underestimates the variation in realised volatility and therefore understates net capital requirements. This motivates the use of scenario generation models such as Stochastic Volatility Jump Diffusion (SVJD) that can be calibrated to produce a wider variation in realised volatility consistent with historical daily returns:

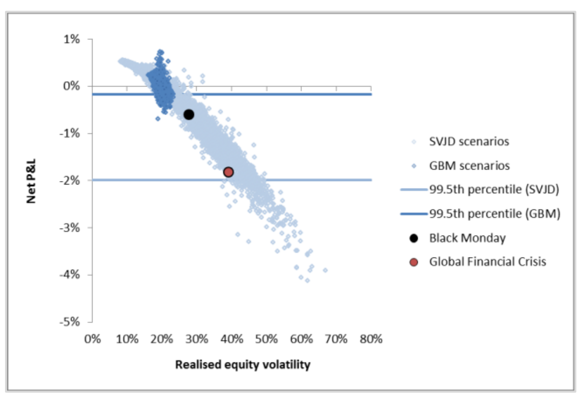

Contents

1. Introduction

In recent years many insurers have invested heavily in the development of internal models for the calculation of 1-year VaR, for example to estimate the regulatory Solvency Capital Requirement as part of Solvency 2.

To date these internal models have typically been designed as *single period* models. Such models generate scenarios for risk factors at a one year horizon (typically by sampling these from a specified joint distribution) and evaluate the resulting asset and liability values at that time (typically using some kind of proxy model). Such an approach assumes that asset and liability values depend only on the values of risk factors at one year. However, in general, the value of assets and liabilities at the one year horizon depend not only on the value of risk factors at the one year horizon but also on their values at intermediate points over the course of the year.

In particular, assets may be regularly rebalanced, often with the specific intention of managing the mismatch between asset and liability values. For example, for certain liabilities, such as guarantees embedded in variable annuity products, dynamic hedging strategies are often in place which attempt to rebalance assets so as to offset the liability's sensitivity to short-term changes in market risk factors such as equity and interest-rate levels. Properly accounting for the effects of such dynamic asset strategies on 1-year VaR capital therefore requires the ability to calculate:

1. Returns on hedging assets over each rebalancing period. This requires simulated risk factor *paths*, not just year one values.
2. The amounts of hedging assets held over each rebalancing period. This requires the evaluation of liability market risk sensitivities (Greeks) at each time-step in each path.

These calculations present new technical challenges to the builders of internal models. In this paper, we explore these technical challenges using a case study based on a stylized Variable Annuity product. The rest of the paper is structured as follows:

- Section 2 describes the first technical challenge: the generation of daily risk factor paths.
- Section 3 describes the second technical challenge: the efficient evaluation of liability Greeks.
- Section 4 presents a case study, demonstrating how a 1-year VaR capital can be calculated for a stylized variable annuity product in the presence of a daily delta hedging strategy.
- Section 5 measures the sensitivity of results to hedging frequency.
- Section 6 concludes.

2. Generation of daily equity paths

In this section, we consider the first technical challenge: the generation of daily paths, focusing on equity returns. What are the key features of historical daily equity returns that are relevant to the assessment of hedge performance, and what kind of scenario generation model can we use to capture these features?

Figure 1 shows historical daily log-returns on the FTSE 100 Total Return Index from 2 January 1986 to 31 October 2014[1]. Note that the amount of variability in returns appears to vary over time, in particular that that periods of extreme returns appear in clusters (for example during the Global Financial Crisis around 2008), and that there are occasional large 'jumps' (for example the 'Black Monday' crash of October 1987).

Figure 1: Historical FTSE 100 daily log-returns (2 January 1986 to 31 October 2014)

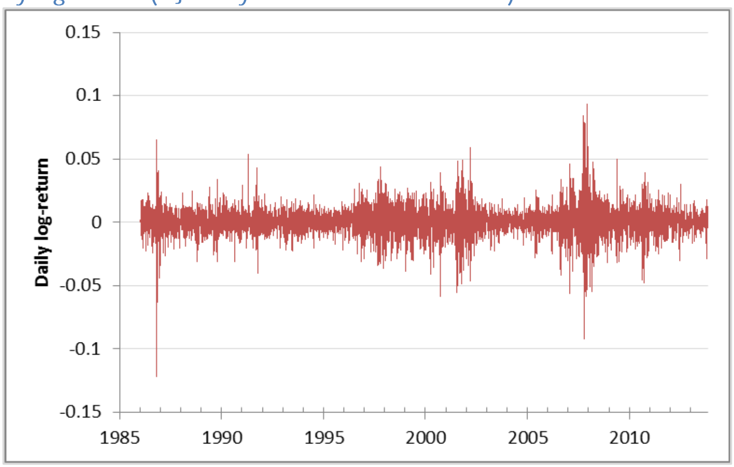

From the point of view of daily delta hedging, relatively large daily returns result in a hedging loss (while relatively low daily returns result in a hedging gain)[2]. Considering the accumulated Profit and Loss (P&L) over a one year horizon, we therefore expect (and will demonstrate later in the paper) that a key driver in hedge performance will be the realised volatility of daily returns over the year.

Figure 2 shows rolling historical realised volatilities (based on one year of daily returns) on the FTSE 100 Index from 31 December 1986 to 31 October 2014. We see a large variation in realised volatility over this period, from a low of 7.9% (year to 6 October 2005) to a high of 39.1% (year to 24 June 2009).

Figure 2: Historical FTSE 100 realised volatilities (31 December 1986 to 31 October 20)

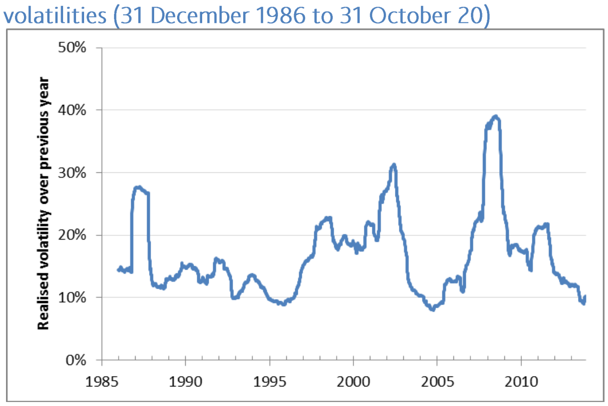

[1] FTSE 100 Total Return Index data provided by Thomson Reuters DataStream.

[2] It is well known - see (Sepp 2012) for example - that the daily P&L of an option position, net of delta hedging and assuming no transaction costs, is approximately given by $\frac{1}{2}S^2\Gamma\delta t(\sigma_{IMPLIED}^2 - \sigma_{REALISED}^2)$ where S is the opening equity price, Γ is the option's gamma, $\sigma_{IMPLIED}$ is the assumed implied volatility (used to both price the option and calculate its delta) and $\sigma_{REALISED} = \frac{1}{\sqrt{\delta t}}\left|\frac{\delta S}{S}\right|$ is the realised annualised volatility (based on a single return observation). Thus hedging losses/gains are made if the realised variance is higher/lower than the implied variance. This approximation is independent of the process used to generate the realised returns.

So which model should we use to generate daily equity returns for the purpose of assessing hedge performance? One candidate model that captures many of the observed features in historical returns (including volatility clustering and jumps) is the SVJD model (Moody's Analytics 2013).

Figure 3 shows the distribution of realised daily UK equity volatility, over a one year horizon, produced by Moody's Analytics standard real-world calibration of the SVJD model for at 31 December 2013. For comparison, we also show the distribution of realised volatility produced by a simpler GBM model[3], and overlay the maximum and minimum historical volatilities of FTSE 100 Index as observed over the period 31 December 1986 to 31 October 2014.

In the context of 1-year VaR capital requirements net of delta hedging, the assumed 99.5th percentile of realised equity volatility is a key assumption. Given the relatively small number of independent historical observations of the realised volatility over one year, there is clearly an element of subjectivity in judging whether the distributions produced by any model provides a good fit to historical data, particularly in the extreme tails. However, the distributions produced by the simple GBM model here appear to be unrealistically narrow given historical experience, while the SVJD calibration produces a far more plausible range, with a 99.5th percentile realised volatility of 45%. For comparison, note that the maximum realised volatility observed for the FTSE 100 (39%, year to 24 June 2009) corresponds to the 98.5th percentile of the SVJD distribution (a 1 in 65 year event)[4]. It is ultimately the judgment of the firm and its regulator whether these are suitable assumptions.

Figure 3: Distribution of realised volatility for GBM and SVJD models (0.5/5/25/50/75/95/99.5 percentiles and average)

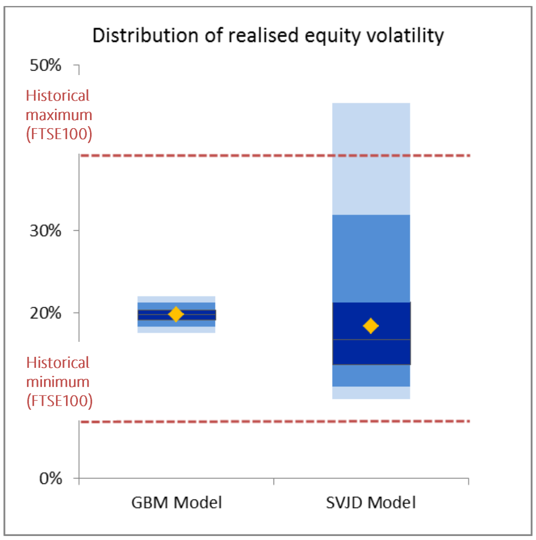

[3] The GBM model has been calibrated to be match the first two moments of the distribution of the annual equity return produced by the SVJD model. Compared to the GBM model, the SVJD model produces excess kurtosis and negative skew in the distribution of the annual equity return.

[4] Note that the SVJD calibration is sensitive to market option implied volatilities at the calibration date (here 31 December 2013) and so can be expected to change according to different market conditions i.e. the distribution in Figure 3 should be interpreted as a conditional ("point-in-time") distribution.

3. Evaluation of liability Greeks using Least Squares Monte Carlo

Having generated daily equity paths, the second technical challenge is the evaluation of liability market risk sensitivities (Greeks), at each daily time-step, in each path.

In most practical hedging programs, the evaluation of Greeks is usually carried out using a 'bump and revalue' approach with each valuation being carried out using risk-neutral stochastic scenarios. A simulation approach to valuation is usually considered necessary due to the complex, path dependent, nature of variable annuity guarantees.

If such an approach were used to project Greeks under a set of real-world stochastic scenarios, this leads to a nested stochastic simulation, as illustrated in Figure 4Figure 1. Here, the red arrows represent the real-world ('outer') scenarios under which the hedge portfolio is projected, while the blue arrows represent the risk-neutral ('inner') scenarios used to value the liability portfolio. The bump-and-revalue approach involves revaluing the liability under a small change to the initial value of the relevant risk factor (for example equity index level in the case of delta) and estimating the resulting sensitivity using a finite difference approximation, so that in practice a number of different valuations are required at each time-step, in each path[5].

Figure 4: Illustration of Greeks projection using nested stochastic method

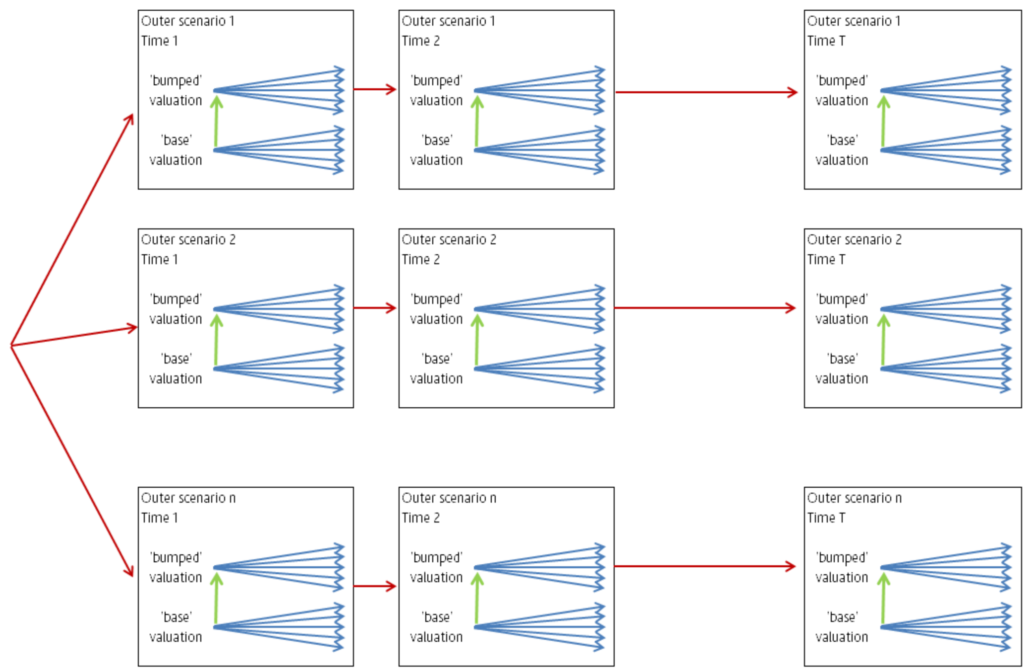

We will see below that many billions of inner scenarios may be required to perform such a nested stochastic simulation, and running such a large number of scenarios through liability cash flow models is usually infeasible given current model run-times. However, many insurers have faced similar nested stochastic simulation challenges already, notably in the projection of 1-year liability values in the calculation of 1-year VaR capital requirements (not accounting for dynamic hedging), and some have already successfully tackled this challenge via the implementation of 'proxy' methods such as Least Squares Monte Carlo (LSMC). In this paper, we consider how these LSMC implementations can be extended to account for dynamic hedging strategies, thus allowing firms to take credit for such hedging strategies in their 1-year VaR capital requirements[6].

[5] Usually a central difference approximation is adopted in practice, with both bumped up and bumped down values being used to estimate sensitivities.

[6] For further background on the use of LSMC in the multi-period projection of Greeks, see previous MA research reports (Clayton, Morrison, Turnbull, & Vysniauskas, Dynamic Hedge Projection and the Multi-Period Modelling of Greeks, 2013) and (Clayton, Morrison, Turnbull, & Vysniauskas, Multi-Period Modeling of Greeks Using Least Squares Monte Carlo: An Exotic Option Case Study, 2014).

LSMC describes the liability value and Greeks by analytical proxy functions as a fast alternative to risk-neutral simulation. Once this function is available, no time-consuming risk-neutral simulation is required. It should be noted that LSMC doesn't remove the need for risk-neutral simulation completely, as a number of risk-neutral scenarios are required to calibrate the proxy functions. However, the number of scenarios that are required to accurately calibrate the function is usually far less than the number required in a full nested stochastic calculation. For illustration, Figure 5 below compares typical scenario requirements in the projection of a daily delta hedging strategy over one year using brute force nested stochastic simulation, compared to the scenarios required to calibrate the associated Greeks using LSMC[7]. LSMC potentially reduces the number of outer scenarios, time-steps and inner scenarios that are required, in this case reducing the total number of inner scenarios required (and hence total runtime) by a factor of almost 100,000.

Figure 5: Scenario requirements for full nested stochastic simulation vs proxy function calibration (LSMC)

	Number required	
	Nested stochastic simulation	Proxy function calibration (LSMC)
Number of outer scenarios	50,000	5,000
Number of time-steps	252	13
Number of valuations per outer scenario/time-step	2	2
Number of inner scenarios per valuation	1,000	2
Total number of inner scenarios (thousands)	25,200,000	260

[7] The example here is illustrative and the exact numbers that are appropriate will vary according to the specific application. Here we assume that a single first order Greek is required, that hedging is carried out on a daily basis (and the model reflects this), that 1,000 scenarios (bumped up and down) are required to accurately estimate each delta and that 50,000 outer scenarios are required to accurately estimate the 99.5th percentile. The LSMC parameters here correspond to those used in the case study in the next section.

4. Variable Annuity case study

To illustrate the impact of a daily delta hedge on 1-year VaR capital requirements, and the sensitivity to the model used to generate daily equity paths, we consider a stylized Variable Annuity product with a Guaranteed Minimum Accumulation Benefit (GMAB). The main product assumptions are as follows:

- 10 year maturity.
- Underlying fund = 30% equities/70% risk-free bonds, rebalanced continuously.
- Guarantee assumed to ratchet continuously with the value of the fund.
- Guarantee charge expressed as fixed % of fund value, paid continuously.

For simplicity we assume there is no mortality or lapses. Under these assumptions, the guarantee in this product is a commonly known as a floating strike look back put option. For valuation purposes we assume a Black-Scholes model with constant risk free rate and equity implied volatility (set equal to the assumed volatility of the one year real-world equity return, 19.8%). Under our assumptions, the t=0 guarantee cost is 13.7% of the initial fund value. The guarantee fee is set so that the t=0 value of future fees exactly matches this guarantee cost i.e. guarantee charge = 1.37% per annum.

Projection of gross P&L

First consider the calculation of 1-year VaR capital requirements without any hedging strategy in place. In order to perform this calculation, we calibrated a polynomial[8] proxy function for the guarantee cost[9] at the one year horizon, using 50,000 outer scenarios and 2 inner scenarios per outer scenario. The 50,000 outer scenarios used to calibrate the proxy function were generated using Moody's Analytics' real-world calibration of the SVJD model at 31 December 2013. The resulting calibrated proxy function is a quadratic function with 6 terms[10]:

$$\begin{aligned}&\mathit{Guarantee\ Cost}\ (\mathit{at\ year\ one})\\ &\quad = 0.0324 + 0.146 \times \mathit{Fund\ Value} - 0.074 \times \mathit{Ratchet} + 2.07 \times \mathit{Fund\ Value}^2 + 2.10 \times \mathit{Ratchet}^2 - 4.14 \times \mathit{Fund\ Value}\\ &\quad \times \mathit{Ratchet}\end{aligned}$$

Having calibrated this proxy function, we need to validate how accurately it approximates the 'true' guarantee cost at the one year horizon, under a range of outer scenarios. In general, the 'true' values of variable annuity guarantees cannot be calculated analytically and so are estimated using simulations (with a relatively large number of risk-neutral scenarios in order to produce an accurate estimate). The resulting runtime constrains the number of outer scenarios that we can validate the proxy function under. However, for the stylized product considered here, and assuming a Black-Scholes valuation model, the value of this look back option can actually be calculated analytically, allowing us to quickly validate the proxy function in a relatively large number of outer scenarios. Indeed, we can validate the proxy function in all scenarios to which the proxy function will be applied (in the current context a set of one year real-world scenarios that will be used to estimate 99.5% VaR).

In Figure 6 below we present validation in 50,000 scenarios produced using Moody's Analytics' real-world calibration of the SVJD model at 31 December 2013[11]. Figure 6 compares the gross P&L (i.e. change in net asset value in the absence of hedging) calculated using the proxy function with that calculated using the analytical valuation formula. Gross P&L is expressed as a percentage of the initial fund value.

[8] For the problem considered here, we find that polynomials provide an accurate description of both the guarantee cost and the delta. In our experience, alternative models (such as Neural Networks) are sometimes required to provide an accurate description of the guarantee cost and Greeks in certain circumstances, in particular close to the guarantee maturity.

[9] The proxy function was calibrated to the cost of future policyholder benefits only. In the calculation of 1-year VaR, the value of guarantee fees was calculated analytically.

[10] All coefficients in this polynomial have been rounded to 3 significant figures for presentation only.

[11] Note that that real-world SVJD scenarios were also used to calibrate the proxy function, but here a different set of 50,000 scenarios have been used for validation purposes – the validation scenarios used can therefore be considered 'out-of-sample' from those used to calibrate the function.

Figure 6: Gross P&L (Proxy vs Analytical) in 50,000 SVJD scenarios

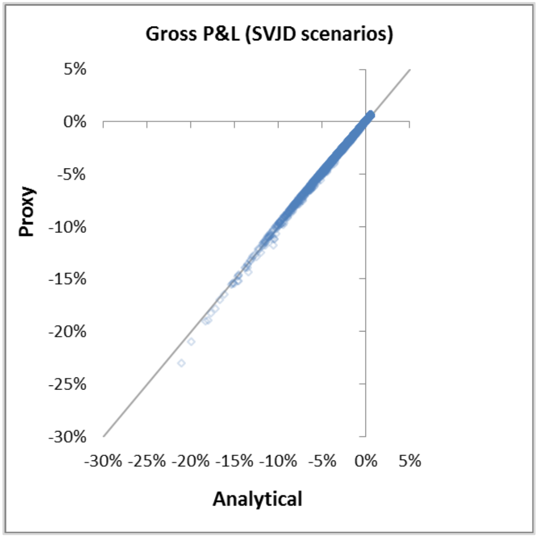

Figure 7 provides further analysis of the gross P&L errors. The leftmost chart plots the error in gross P&L (proxy – actual) against the annual equity return. This indicates that the largest errors tend to occur at the most extreme equity returns. This is expected due to the lower density of outer scenarios used to calibrate the guarantee cost proxy function at these extreme values[12]. The middle chart summarizes the distribution of errors, indicating that the errors are generally small relative to the initial fund value (0.5^{th} and 99.5^{th} percentiles are -0.12% and 0.15% respectively). In order to get a further sense of scale, the rightmost chart expands the y-axis and adds two key balance sheet items for comparison - the t=0 guarantee cost (13.7%) and the gross capital requirement (6.61%). Errors in gross P&L are generally small relative to both of these quantities.

Figure 7: Analysis of Gross P&L errors (Proxy - Analytical) in 50,000 SVJD scenarios

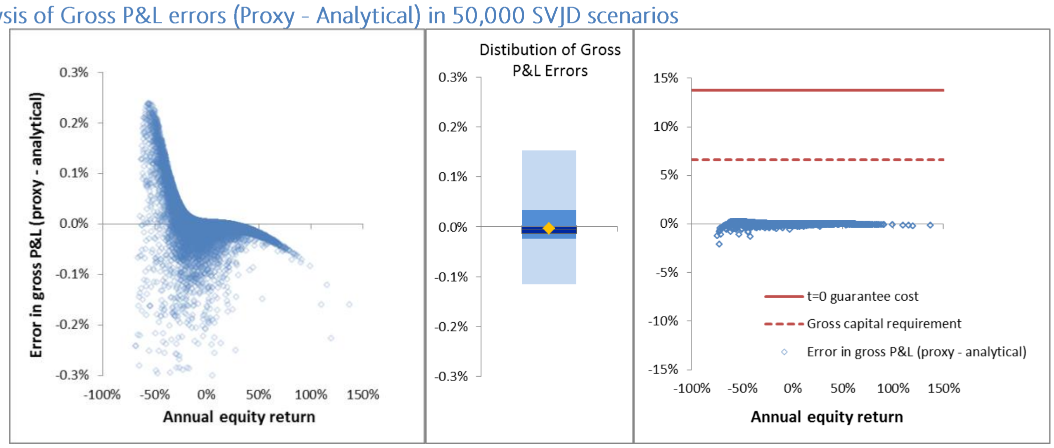

In summary, the proxy function for the guarantee cost appears to provide a good approximation to the true guarantee cost in the SVJD scenarios. Corresponding validation results using a set of scenarios produced using an alternative GBM model are provided as an Appendix to this note, with similar conclusions.

[12] Recall that these outer scenarios used to calibrate the proxy function were generated using the SVJD model. We could use alternative outer scenarios in the calibration and in particular place higher density at the more extreme equity returns in an effort to improve the fit in these regions if required.

Having validated the guarantee cost proxy function, we can proceed to estimate the gross 1-year VaR capital requirement. Figure 8 plots gross P&L (as estimated using the guarantee cost proxy function) against one year equity return in 50,000 SVJD scenarios and 50,000 GBM scenarios, with the 99.5th percentile losses shown for both choices of real-world model.

Figure 8: Gross P&L vs annual equity return – SVJD model (LHS); GBM model (RHS)

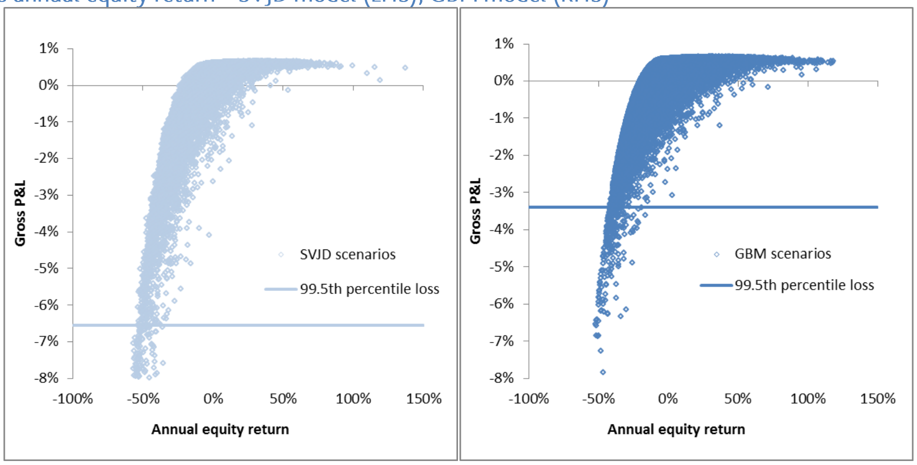

As expected, the gross capital requirement is far higher using the SVJD model due to its fatter downside equity return tail (the SVJD and GBM models here are calibrated to have the same mean and standard deviation, but the SVJD model has excess kurtosis and negative skew).

Projection of hedge portfolio value

Now consider the addition of a daily delta hedging strategy. To enable simulation of such a strategy, we additionally calibrated a second polynomial proxy function, in this case for the guarantee delta[13]. While the guarantee cost only needs to be evaluated at the one year horizon, the delta needs to be evaluated daily. We therefore include time as an additional variable in the delta proxy function. To calibrate the delta proxy function, we have used 5,000 outer scenarios and 2 inner scenarios per outer scenario, at monthly time-steps from time zero to one year i.e. 13 time-steps in total (recall Figure 5). As before, the 5,000 outer scenarios used to calibrate the delta proxy function were generated using Moody's Analytics' real-world calibration of the SVJD model at 31 December 2013. Delta was estimated using a central difference bump-and-revalue estimator. The resulting calibrated proxy function is a cubic function with 9 terms[14]:

$$Guarantee\ Delta = -24.3 - 0.00972 \times Time + 17.6 \times Fund\ Value + 51.6 \times Ratchet + 40.6 \times Fund\ Value^2 - 5.10 \times Ratchet^2$$
$$-101 \times Fund\ Value \times Ratchet - 35.6 \times Fund\ Value^3 + 56.0 \times Fund\ Value^2 \times Ratchet$$

Note that this guarantee delta proxy function explicitly depends on time, via the second term (linear in time), though this time dependence is relatively small in magnitude. This lack of explicit time dependence reflects the relatively long maturity of the guarantee (10 years at t=0) relative to the projection horizon (one year).

[13] As before, the proxy function was calibrated to the delta of future policyholder benefits only and the delta of guarantee fees was calculated analytically.

[14] All coefficients in this polynomial have been rounded to 3 significant figures for presentation only.

As before, validation is aided by the availability of analytical functions for the delta[15]. Rather than validating the quality of fit of the delta itself, at each time-step, here we focus our attention on the monetary impact of using this delta i.e. the resulting P&L of the hedge portfolio at the one year horizon.

Figure 9 compares the hedge portfolio P&L calculated using the proxy function with that calculated using the analytical formula, in 50,000 SVJD scenarios, while Figure 10 provides further analysis of the hedge portfolio P&L errors. As for the gross P&L, the largest errors tend to occur at the most extreme equity returns, as expected given the lower density of outer scenarios used to calibrate the guarantee delta proxy function at these extreme values. In general these errors are generally small relative to initial fund value, t=0 guarantee costs and the gross capital requirements.

Figure 9: Hedge portfolio P&L (Proxy vs Analytical) in 50,000 SVJD scenarios

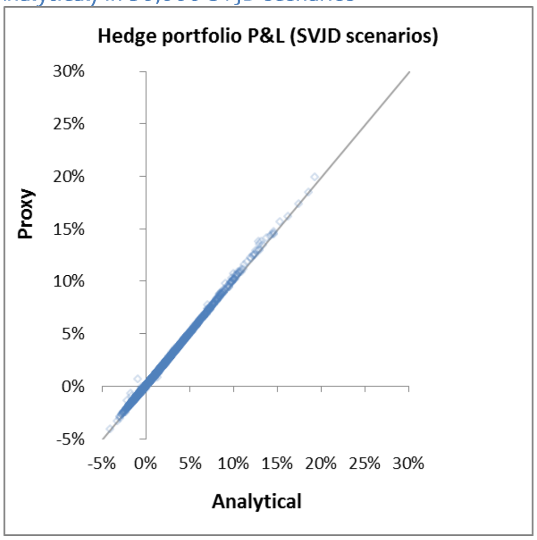

Figure 10: Analysis of hedge portfolio P&L errors (proxy - analytical) in 50,000 SVJD scenarios

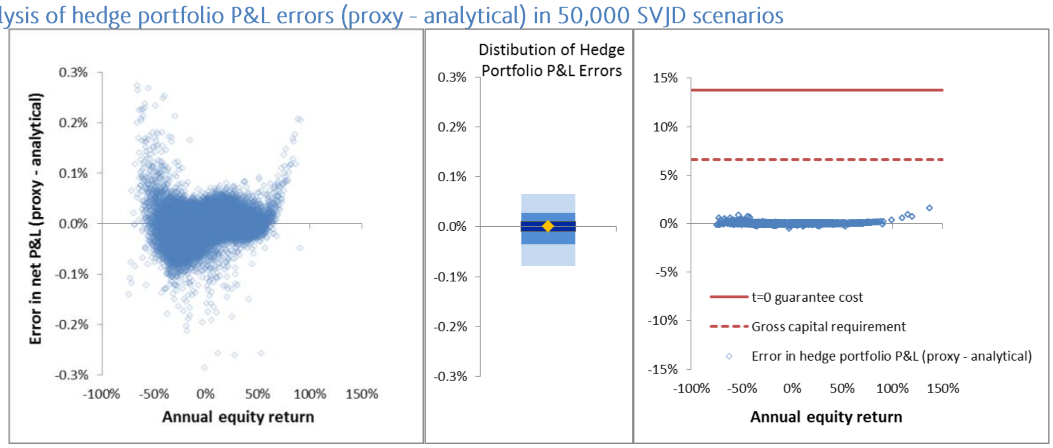

[15] In practice, the analytical deltas here were calculated using a central difference approximation, with each valuation calculated using analytical valuation formulae.

The validations in figure 9 and 10 suggest that the guarantee cost and delta proxy functions result in accurate values for the gross and hedge portfolio P&Ls respectively, across 50,000 SVJD scenarios. However, the key variable of interest in the calculation of net capital requirements (in the presence of the delta hedge), is the net P&L i.e. the sum of the gross and hedge portfolio P&Ls.

Putting it all together: Net P&L and capital requirements

Figure 11 compares the net P&L calculated using the proxy function with that calculated using the analytical formula, in 50,000 SVJD scenarios, while Figure 12 provides further analysis of the net P&L errors. The distribution of net P&L errors is only slightly wider than the distribution of gross P&L errors - in absolute terms, the net P&L errors are of similar size to the gross P&L errors. In relative terms, these errors are larger for net P&L, as the true distribution of net P&L is narrower than the distribution of gross P&L (reflective the effectiveness of the hedge). Nonetheless, the errors are still small relative to the net capital requirement (1.99%).

Figure 11: Net P&L (Proxy vs Analytical) in 50,000 SVJD scenarios

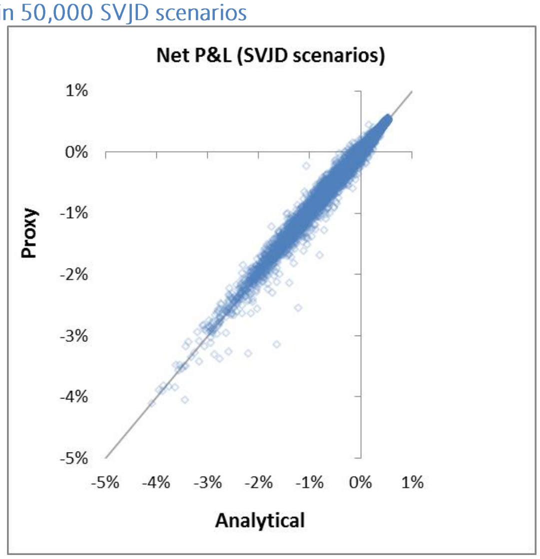

Figure 12: Analysis of net P&L errors (proxy - analytical) in 50,000 SVJD scenarios

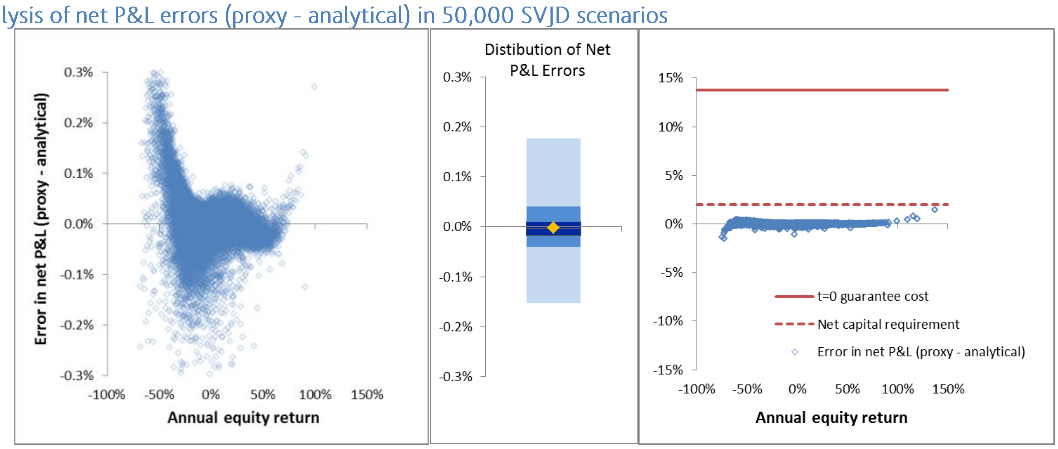

Corresponding validation results using a set of scenarios produced using an alternative GBM model are provided as an Appendix to this note. The major difference in using GBM scenarios is that the resulting net capital requirement is close to zero in this case i.e. the GBM model implies an almost perfect delta hedge, so that the net P&L errors are larger relative to the net capital requirement (though still small relative to the other balance sheet items).

Having validated the guarantee cost and delta proxy functions, we can proceed to use them to estimate the net 1-year VaR capital requirement. Figure 13 plots net P&L (as estimated using the guarantee cost and delta proxy functions) against realised equity volatility in 50,000 SVJD scenarios and 50,000 GBM scenarios, with the 99.5th percentile losses shown for both choices of real-world model. In addition to the stochastic scenarios, two historical scenarios with relatively large realised equity (FTSE 100) volatilities are also shown for comparison:

- ‘Black Monday’ (year to 28 March 1988), a year that included two consecutive daily FTSE 100 falls of over 10% with a realised volatility of 28%
- ‘Global Financial Crisis’ (year to 24 June 2009), the year with the highest realised FTSE 100 volatility to date (39%)

Figure 13 confirms that realised equity volatility is the key driver of net P&L, with net losses tending to occur in scenarios that have realised volatility in excess of the implied volatility (recall the discussion of Section 2). The unrealistically narrow distribution of realised volatility produced by the GBM model results in a net 1-year VaR capital requirement that is far lower than the corresponding loss made in the ‘Black Monday’ and ‘Global Financial Crisis’ scenarios. In contrast, the wider distribution of realised volatility produced by the SVJD model results in a net 1-year VaR capital requirement just larger than the loss made in the Global Financial Crisis scenario (which, from the delta hedger’s point of view, can be considered the worst historical scenario to date).

Figure 13: Net P&L vs realised equity volatility in 50,000 stochastic scenarios (SVJD and GBM models) and 2 historical scenarios (Black Monday and Global Financial Crisis)

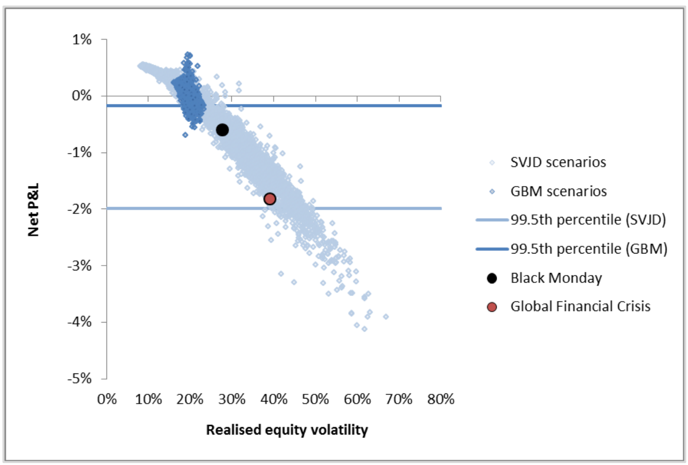

To summarize, Figure 14 shows the 1-year VaR capital requirements implied by the two different real-world models, along with a measure of ‘hedge effectiveness’ (defined as $1-\frac{\text{Net 99.5\% VaR}}{\text{Gross 99.5\% VaR}}$).

Figure 14: Gross and net capital requirements estimated using proxy functions

	99.5% VaR		Hedge
Real-world scenario generator	Gross	Net	effectiveness
Geometric Brownian Motion	3.39%	0.17%	95%
SVJD	6.56%	1.99%	70%

5. Sensitivity to hedging frequency

In our assessment of hedge effectiveness above, we have simulated a delta hedging strategy at daily frequency. Though this reflects the frequency of delta hedging often used in practice, from a modelling point of view it is natural to ask how the hedge effectiveness changes if we assume the hedge is rebalanced less frequently. Assuming a lower rebalancing frequency within the model (such as quarterly) may only approximate the true hedging strategy but can be expected to provide at least some reduction in capital requirements while potentially being computationally feasible using a full nested stochastic model (since the number of delta evaluations required is significantly reduced compared to a daily hedge). As an extreme case, an assumption that the hedge is never rebalanced over the course of the year only requires the delta to be calculated once (at t=0).

If we assume a reduced hedging frequency, how effective is the resulting hedge in reducing 1-year VaR capital requirements? Figure 15 shows the distribution of net P&L at different hedging frequencies, and compares with the gross (unhedged) P&L. In all cases, the SVJD model has been used to generate equity scenarios.

While a daily hedge results in a 70% reduction in capital, this drops to 44% at quarterly rebalancing frequency, and only 2% at annual frequency. At annual frequency, the hedge has virtually no effect.

Figure 15: Distribution of P&L at various hedging frequencies (0.5/5/25/50/75/95/99.5 percentiles and average)

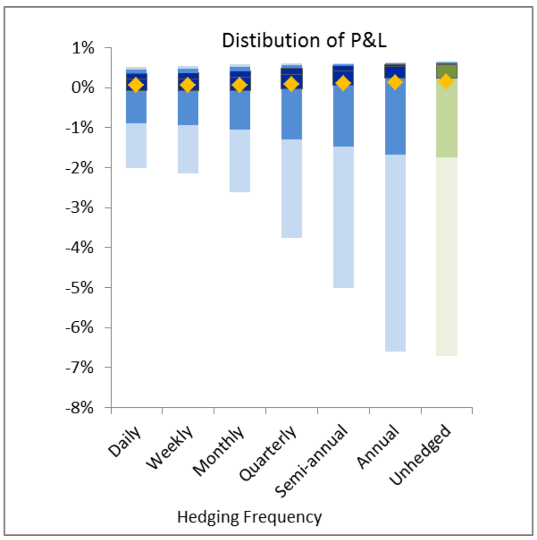

To understand this result, Figure 16 shows the value of guarantee benefits and fees (at t=0), as a function of the underlying fund value, while Figure 17 shows the corresponding deltas[16]. When the fund value is 1, the value of fees + benefits is exactly zero by construction – the annual guarantee fee has been set so that the value of fees exactly matches the value of benefits. More interestingly, the guarantee price as a function of fund value flattens off at this point, so that the guarantee delta is close to zero. As a result, there is virtually no hedge in place at t=0 and an assumption of no rebalancing is almost equivalent to an assumption of no hedging. The shape of the delta is however highly convex at this point and the delta quickly increases as the fund value falls. Thus for the product considered here, an effective hedge strategy requires relatively frequent hedging. It should be noted that these results are sensitive to the features of the product, in particular the initial moneyness of the guarantee.

[16] The ratchet is assumed to be the larger of 1 and the underlying fund value.

Figure 16: t=0 guarantee value vs underlying fund value

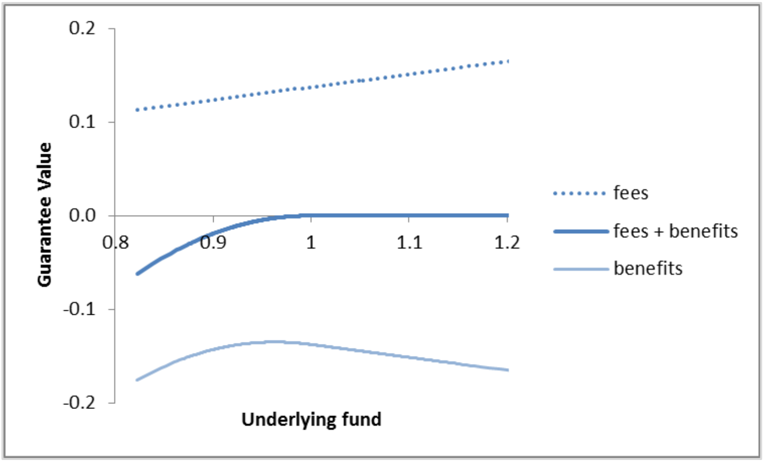

Figure 17: t=0 guarantee delta vs underlying fund value

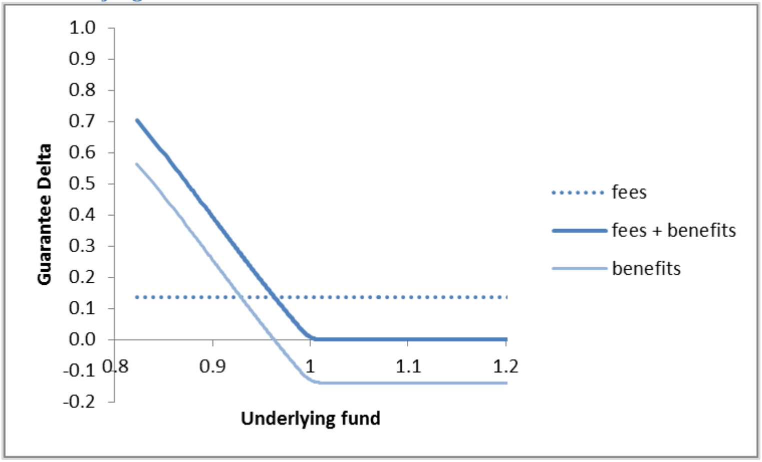

6. Summary

In this note, we have considered some of the technical challenges and solutions in adapting internal models to account for the effect of dynamic hedging strategies in reducing 1-year VaR capital requirements.

The 'single period' models commonly adopted to assess 1-year VaR are designed to capture risk at a single (one year) time horizon and therefore do not naturally extend to include the effects of dynamic asset (or liability) management strategies over the year. Capturing the effect of dynamic hedging raises new technical challenges to the builders of internal models, in particular:

1. Generation of economic scenarios for realistic daily risk factor paths.
2. Fast evaluation of liability Greeks.

We have illustrated how these challenges can be tackled using a case study, based on a variable annuity business with a daily delta hedge. Though focused on equity risk and delta hedging, the techniques described in this paper readily extend to other key risks, in particular interest rates. Liability Greeks are approximated using accurate proxy function techniques, calibrated using LSMC.

This case study highlights the sensitivity of hedge performance to the realised daily volatility of underlying risk factors, and the resulting sensitivity of 1-year VaR capital requirements to the assumed variation in volatility. A relatively simple GBM model severely underestimates the variation in realised volatility and therefore understates net capital requirements. This motivates the use of scenario generation models such as SVJD, that can be calibrated to produce a wider variation in realised volatility consistent with historical daily returns.

7. Appendix

In Section 4, we considered validation of value and delta proxy functions under a set of stochastic scenarios generated using the SVJD model. Here we consider validation under a set of 50,000 stochastic scenarios generated using the GBM model.

Firstly, we consider validation of the proxy function for the time one guarantee cost. Figure 18 compares the gross P&L calculated using the proxy function with that calculated using the analytical formula, in 50,000 GBM scenarios, while Figure 12 provides further analysis of the gross P&L errors.

Figure 18: Gross P&L (Proxy vs Analytical) in 50,000 GBM scenarios

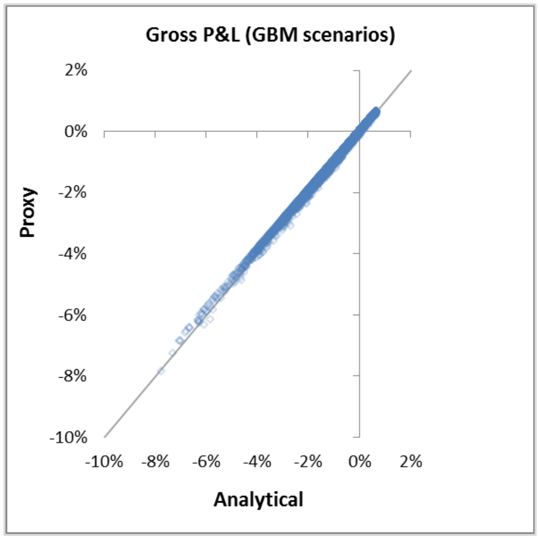

Figure 19: Analysis of Gross P&L errors (Proxy - Analytical) in 50,000 GBM scenarios

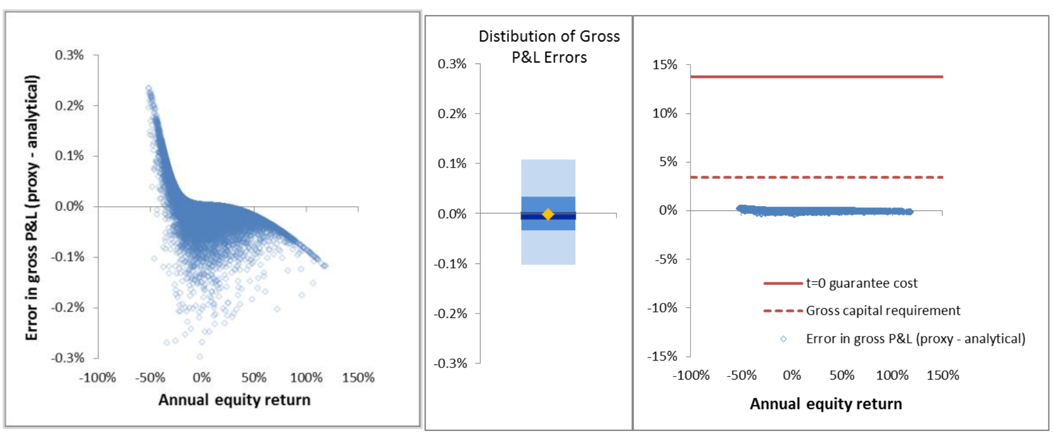

The distribution of errors here is similar to that observed using the SVJD scenarios (recall Figure 7). One difference is that the gross capital requirement under the GBM model (3.39%) is lower than that under the SVJD model (6.56%)[17] and so the errors using the GBM model are larger relative to the corresponding capital requirement than the corresponding relative errors using the SVJD model. Nevertheless, errors are still relatively low compared to the gross capital requirement (and the t=0 guarantee cost).

[17] Both capital requirements here are estimated using the proxy functions.

To validate performance of the proxy function for the guarantee delta, Figure 20 compares the hedge portfolio P&L calculated using the proxy function with that calculated using the analytical formula, in 50,000 GBM scenarios, while Figure 21 provides further analysis of the hedge portfolio P&L errors.

Figure 20: Hedge portfolio P&L (Proxy vs Analytical) in 50,000 GBM scenarios

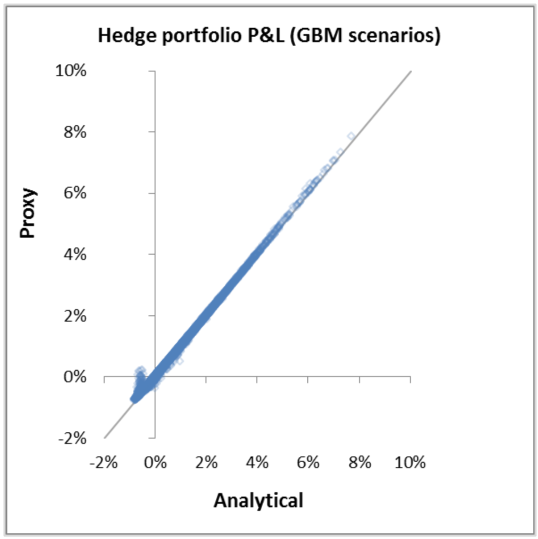

Figure 21: Analysis of hedge portfolio P&L errors (Proxy - Analytical) in 50,000 GBM scenarios

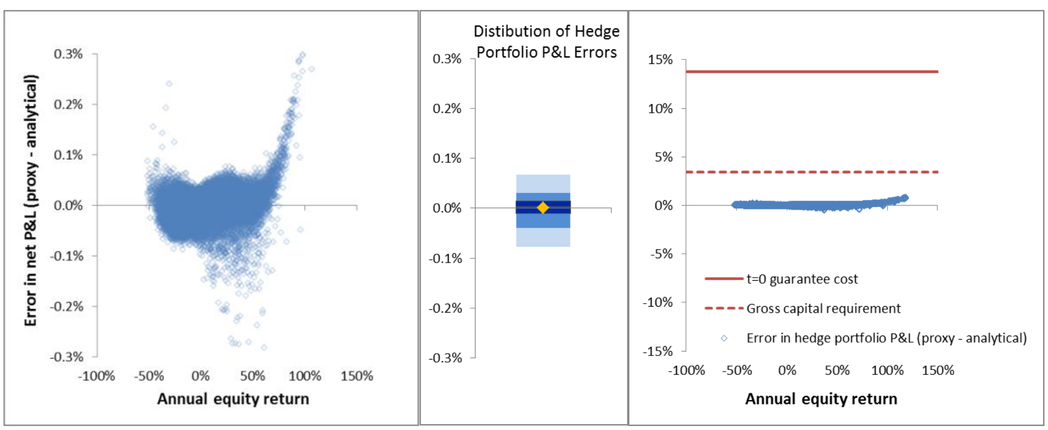

The distribution of errors here is similar to that observed using the SVJD scenarios (recall Figure 10), with the largest errors occurring at the largest equity returns. It is worth noting that the GBM calibration considered here gives rise to more extreme scenarios (with relatively large net P&L errors) on the upside (and less extreme scenarios on the downside) than the SVJD model, due to negative skew in the SVJD model.

Finally, Figure 22 compares the net P&L calculated using the proxy function with that calculated using the analytical formula, in 50,000 GBM scenarios, while
Figure 23 provides further analysis of the net P&L errors.

Figure 22: Net P&L (Proxy vs Analytical) in 50,000 GBM scenarios

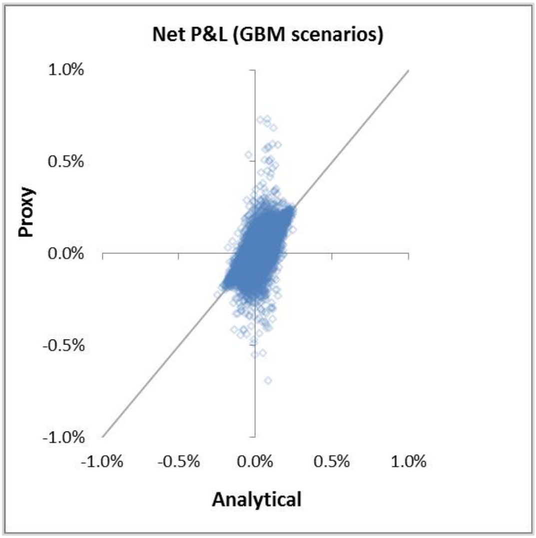

Figure 23: Analysis of net P&L errors (Proxy - Analytical) in 50,000 GBM scenarios

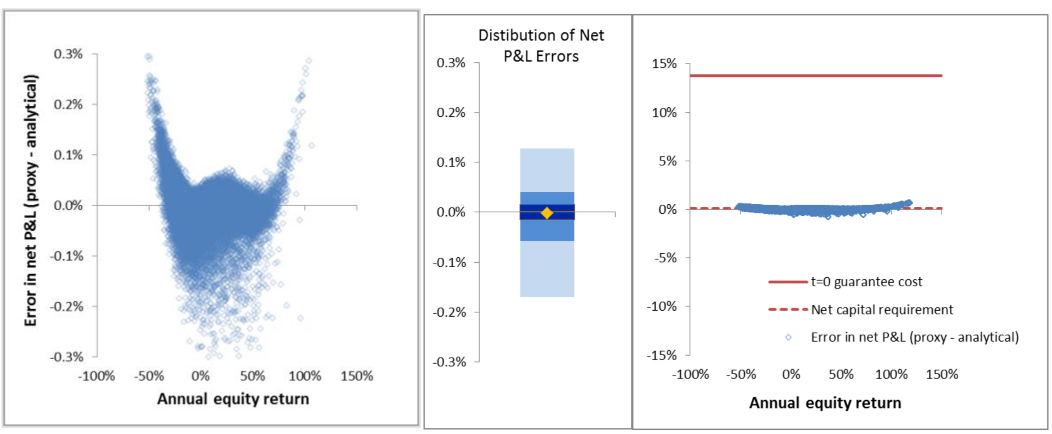

The distribution of errors here is again similar to that observed using the SVJD scenarios (recall Figure 12), with the largest errors occurring at the largest equity returns. One difference is that the net capital requirement under the GBM model (0.17%) is far lower than that under the SVJD model (1.99%)[18] and so the errors using the GBM model are far larger relative to the corresponding capital requirement than the corresponding relative errors using the SVJD model. In fact, the net capital requirement produced using the analytical formulae is 0.12% under the GBM model (compared to 0.17% estimated using the proxy formulae). This may be considered a large error in relative terms, though it is worth noting that it is very small relative to the t=0 guarantee cost (and the initial fund value).

[18] Both capital requirements here are estimated using the proxy functions.

References

Clayton, Aubrey, Steven Morrison, Craig Turnbull, and Naglis Vysniauskas. *Dynamic Hedge Projection and the Multi-Period Modelling of Greeks.* Moody's Analytics Research Report, 2013.

Clayton, Aubrey, Steven Morrison, Craig Turnbull, and Naglis Vysniauskas. *Multi-Period Modeling of Greeks Using Least Squares Monte Carlo: An Exotic Option Case Study.* Moody's Analytics Research Report, 2014.

Moody's Analytics. "Stochastic Volatility Jump Diffusion: Model Definition." 2013.

Sepp, A. "An approximate Distribution of Delta-Hedging Errors in a Jump-Diffusion Model with Discrete Trading and Transaction Costs." *Quantitative Finance*, 2012: 3(1), 19-59.