On the effectiveness of Hedging Strategies for Variable Annuities

Using a simple example of a Put Option this paper explores such matters as the significance of Realised Volatility risk, Underlying Volatility risk and choice of hedged formula on the effectiveness of applying a hedging strategy to Variable Annuity contracts. It also briefly considers the stress tests proposed by the Central Bank of Ireland and Solvency II in this simplified context.

Brian Woods FSAI

Presented to Society of Actuaries in Ireland, April 2014

Contents

Appendixes

Figures

Tables

Introduction

In the idealised model of Black Scholes we know that the first order delta hedge is perfectly effective in the sense that it removes all variability of the final outcome of a derivative product such as an option. The main purpose of this paper is to examine, using a simple example, the extent to which deviations from the idealised framework render practical hedging strategies less than perfectly effective in so much as they present some variability in final outcomes. I also briefly address other aspects in this simplified context such as Solvency II and the recently published CBI stress tests[1] for Variable Annuity (VA) products.

Throughout this paper the example chosen is a 5 year single premium Put option on a stock/index/fund with reinvested dividends. Unless otherwise stated the initial stock price and the exercise price is taken to be 100. This is broadly analogous to the VA GMAB. I will not be considering demographic or other behavioural complications which would be typically present in such a product. Nonetheless I believe that many of the observations are at least directionally informative in a real VA context though practitioners might want to check to what extent they are applicable to their own situation.

The framework for the analysis is, unless otherwise stated, the same 1,000 stochastic simulations of 60 monthly movements in the underlying stock price using the Brownian log-normal with Real World mean of 5% per annum and volatility of 20% per annum. The risk free rate is assumed to be 2% per annum. Hedging strategies are assumed to rebalance at the monthly simulated price change.

This simple example would in reality have two stochastic drivers viz. the stock price level and the interest rate level. However, as there is no optionality present on the interest rate level "hedging" the interest rate is more a matter of duration matching and would tend to be almost 100% effective unless the liability is of a very long duration. As I wish to focus on the effectiveness of hedging strategies in reducing the variability of outcomes where there is optionality in the liability I will assume that the interest rate level is constant.

The effectiveness of a hedging strategy would be quantified by some measure of the dispersion of the result which it is attempting to stabilise. I have in the main used the CTE90[2] of the outcomes as the measure of dispersion as this is the target of most concern. Clearly it would be highly desirable to have a strategy which reduces the worst 10% of outcomes without a broadly symmetrical effect on the best outcomes. Unfortunately that would be in the realm of alchemy. In statistical terms hedging reduces dispersion whilst having a minimal though negative (due to costs) effect on location. Also the unhedged outcome will be broadly symmetrically distributed either side of the hedged outcome.

Where possible I use closed form formulae in calculating the "greeks" and the option values and these are given in Appendix 2.

[1] Stress Testing Framework for the Variable Annuity Industry

[2] The conditional tail expectation, being the average value of the result conditioned on being in the "worst" 10% of outcomes.

The Taylor Expansion

The Taylor expansion of the option value to first order in δt is as follows:

$$\delta Lov = \Delta\delta s + \Theta\delta t + \frac{\Gamma}{2}(\delta s)^2 + o(\delta t) \qquad (1)$$

where

δLov is the small change in the Liability Option Value[3] caused by the evolution of a small interval of time δt which also causes a small change δs in the stock price.

$\Delta = \frac{\partial Lov}{\partial s}$; $\Theta = \frac{\partial Lov}{\partial t}$; $\Gamma = \frac{\partial^2 Lov}{\partial s^2}$ are the said greeks: delta, theta and gamma

$o(\delta t)$ is a residual term which is small relative to δt

Note that in a practical VA situation none of these terms can be evaluated by closed formulae. In practice the liability is simulated by Monte Carlo methods based on an input Economic Scenario Generator. Nonetheless, provided sufficiently many runs are performed the law of large numbers should give us comfort that the Taylor expansion is a good approximation and indeed this is the ubiquitous tool in VA hedging math.

The interesting term in equation (1) is the term in $(\delta s)^2$ but this is the essence of the stochastic nature of stock price movements. δs is of order $(\delta t)^\alpha$ where $\alpha = \frac{1}{2}$; if α was greater than $\frac{1}{2}$ then the change in the stock price would have no variance at the continuous level i.e. it would be deterministic; if it is less than $\frac{1}{2}$ then the process explodes. $\alpha = \frac{1}{2}$ is the only value which produces the random walk which we intuitively require. Of course the geometric Brownian motion satisfies this condition.

Now equation (1) is not really the Black Scholes SDE and nor does it require the elaborate apparatus of the risk neutral martingale theory or even a principle of no arbitrage. It is fairly elementary mathematical common sense with a sprinkling of Itô. It leads in a natural way to the notion that if we want to be hedged against changes in the stock price we should purchase a position in the primary market which precisely cancels the Δ term, the so called delta hedge.

In the case of our Put option that means having a short position in stock futures. But equation (1) would suggest that we would still be exposed to the stochastic term in $(\delta s)^2$ (the sum of these terms over an interval will hereafter be referred as the *realised volatility*[4]). Now by its nature the primary market does not have instruments which can counter this term, or in the jargon they do not have any convexity (gamma). One of the implications of the Black Scholes derivation is that we do not need such instruments. The delta hedge is perfect in the sense that if we follow it we will have no variability in final outcome – it is a complete replication strategy.

To understand what is happening here imagine the unit time interval divided into $1/\delta t$ sub-intervals of size δt. The realised volatility of each sub-interval will have an expectation of order δt and a variance of order $(\delta t)^2$. The realised volatility over the full interval will have an expectation of the order of $\left(\frac{1}{\delta t}\right) \times \delta t$ i.e. non zero but bounded just as we would expect of a random walk. On the assumption that each sub-interval is independently distributed[5] it will have a variance of the order of

[3] LOV is the standard jargon used by VA writers

[4] Realised volatility is slightly different from realised variance in that the latter has an adjustment for the square of the average δs. At the continuous level the two concepts are the same. Also note that I interchangeably refer to volatility as either being the sum of squares or alternatively the square root of that sum. I hope the context will remove any possible ambiguity.

[5] But not necessarily identically distributed

$\left(\frac{1}{\delta t}\right) \times (\delta t)^2$ i.e. of orderδt. In other words, in the continuous limit the realised volatility over any time interval has zero variance i.e. it is deterministic and is precisely the underlying volatility assumed in the Black Scholes derivation.

Note that this convexity or gamma term in the realised volatility always produces a positive increase in the liability, a fact which follows from the curvature of the Black Scholes formula. Appendix 1 gives what I believe to be a more heuristic illustration of why this is the case. The point is that these convexity effects are instantaneously deterministic in the Black Scholes model and built into the value of the option. The delta hedge works in Black Scholes because the theta of the formula i.e. the instantaneous change in the value with the passage of time alone[6] precisely counters these convexity effects. Gamma and theta themselves are stochastic but they change in a self-cancelling way.

So we see that the first practical difficulty we have with a delta hedge is that we cannot trade at the continuous level and therefore we are exposed to stochastic movements in the realised volatility. These can give rise to profits or losses depending on how they compare with the underlying volatility in the formula.

A perhaps even more obvious "risk" that we run with a delta hedge is that we might not be using the correct assumption for underlying volatility. This will generally be estimated from historic data or maybe from the implied volatility of derivative instruments in the market but clearly we are at risk of having this key assumption in our hedging formula wrong. Note that this is the risk of having the wrong delta; it is not in that sense a gamma risk. We are following a delta hedge so the only thing we can get "wrong" is to calculate the wrong delta and we can take some comfort that this will give a broadly symmetrical exposure to stock price movements . Before examining these Realised Volatility and Underlying Volatility risks in our simple model it is useful at this stage to produce a table of the various versions of "volatility" with which we are confronted.

Table 1: The various versions of volatility

Definition	Comments
Underlying	Can never be observed even with hindsight. Has to be estimated from historic data possibly with a judgement overlay. Implied Volatility could be interpreted as the market's assessment of Underlying Volatility.
Realised	What we observe with hindsight. A random realisation of the Underlying Volatility and plausibly a likelihood estimate of it. Depends on the interval of observation. Unaffected by choice of Accounting or Hedge volatilities.
Accounting	Does not affect final outcome. Does affect timing of P&L recognition and in particular the starting value of the option and affects the attribution of P&L between delta and gamma effects.
Hedging	The volatility assumption in the formula we chose to hedge. Only affects Delta. Delta mismatches with the other definitions are broadly symmetrical in their effect.

[6] That is theta is a partial derivative with respect to time

Realised Volatility risk

Figure 1: Variation in hedged P&L due to Realised Volatility

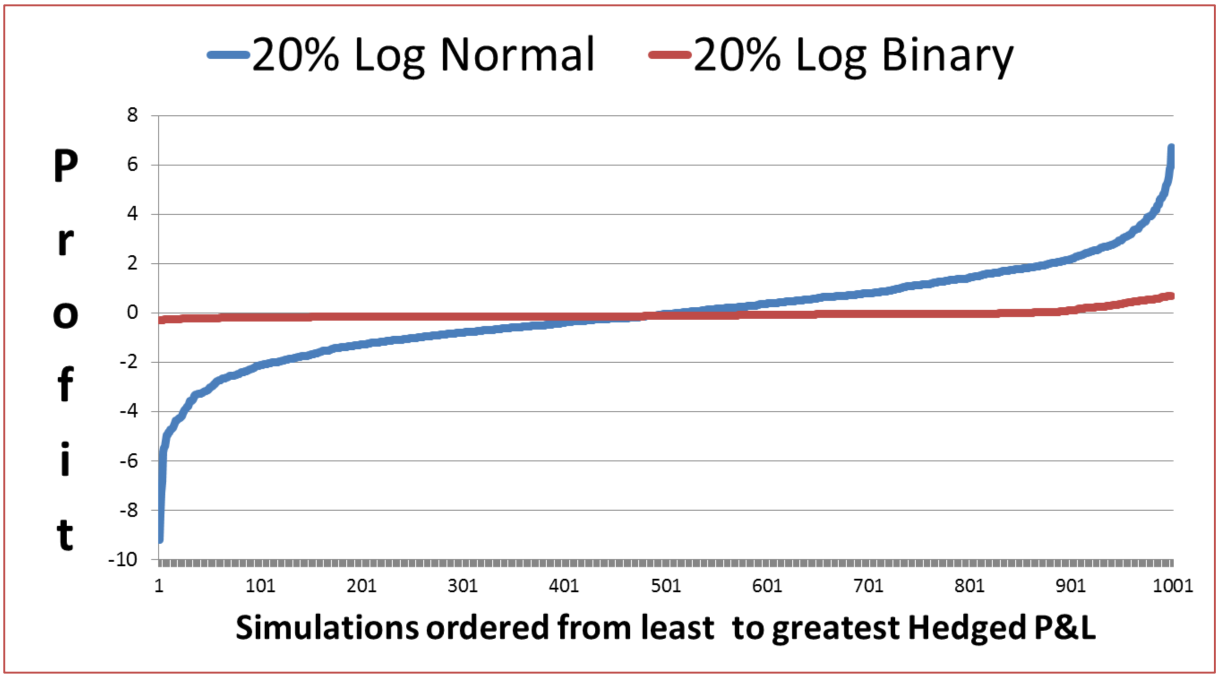

In Figure 1 I consider two possible models for the stochastic behaviour of the stock price. In each case the underlying volatility is 20% and the hedging strategy calculates delta using the Black Scholes log-normal formula with 20% volatility i.e. we are not exposed to Underlying Volatility risk.

In one case I sample from the log-normal distribution as in all other situations in this paper. In the other case I sample from the log-binary distribution. By this latter I mean that each month the log of the price goes up or down by one sigma[7] with equal probability. In this way the realised volatility will be precisely 20% just as it is in the Black Scholes continuous model. So, perhaps paradoxically, the log-binary model more closely represents the outcome from the Black Scholes model than does the log-normal from which the formula is derived and which in each case underpins the calculation of delta. The thought may occur that if we use a log-binary model rather than the log-normal would we derive the same Black Scholes formula for the value of the option? The answer is yes for the log-binary becomes the log-normal in the continuous limit by virtue of the central limit theorem.

The fact that the log-normal outcomes in the above figure are not a horizontal line along the X axis illustrates that there is a residual stochastic dimension to the outcome of the hedging strategy which in this example is driven purely by the Realised Volatility.

[7] Plus a "drift" component

Table 2 shows the output from the stochastic runs using 20% for both the underlying volatility and the hedge volatility.

Table 2: Realised Volatility risk when hedging volatility assumption is correct

	Average Final Outcome (Profit)	Standard Deviation of Final Outcome	CTE90 of Final Outcome
Unhedged	4.5	15.8	-33.8
Hedged using 20% Volatility	0.0	1.9	-3.4

Note: the starting position is taken to be the Risk Neutral value of the Put option. CTE90 is the average of the worst 10% of final profit outcomes.

The Risk Neutral value of the Put option at inception is 12.5 which rolled up to maturity at risk free rate is 13.8. To a first decimal place this resulted in the hedging strategy breaking even on average which reflects the fact that it was the Risk Neutral value of the Put option which was hedged. For the rest of this paper I will refer to this hedging strategy as the Risk Neutral or IFRS strategy. It is the one usually adopted in practice as the IFRS value of the option is determined as its risk neutral expectation value.

The average outcome of the hedged strategy is somewhat worse than the average position on the unhedged strategy since in this latter case the Real World equity risk premium of 3% p.a. reduces the extent of the Put finishing in the money or, alternatively, any hedging strategy in this example will be shorting the equity risk premium.

The hedge dramatically reduces the dispersion of final outcomes both in terms of standard deviation and CTE90. The CTE90 hedge effectiveness can be calculated as the percentage reduction from the unhedged value of -33.8 to -3.4 which is c.90%.

Nonetheless the main point of Table 2 is to show that there is a residual variability, albeit relatively small, notwithstanding that the formula which was hedged precisely described the underlying reality in terms of volatility i.e. by construction we are assured that there is no Underlying Volatility risk.

Of course the Realised Volatility risk is a function of the size of the interval for rebalancing. The underlying model is that realised volatility i.e. the realised square of the price changes is of the order of the time interval and its variance is of the order of the square of the time interval. So if we split the time interval into n sub-intervals each sub-interval will display a variance in the square of the price change of an order $1/n^2$ the variance for the whole interval. Summing these independent variables over the sub-intervals we find that the variance of the realised volatility over the whole interval is 1/n times the variance over the whole interval without subdivision. Hence the Realised Volatility risk reduces in proportion to the square root of the size of the re-balancing interval. For example the Realised Volatility risk for weekly rebalancing might be only 50% of that for monthly rebalancing. The choice of frequency of rebalancing is largely a trade-off between hedge effectiveness and transaction costs.

Before leaving this section we consider Figure 2 below.

Figure 2: Hedging P&L outcome plotted against final stock price

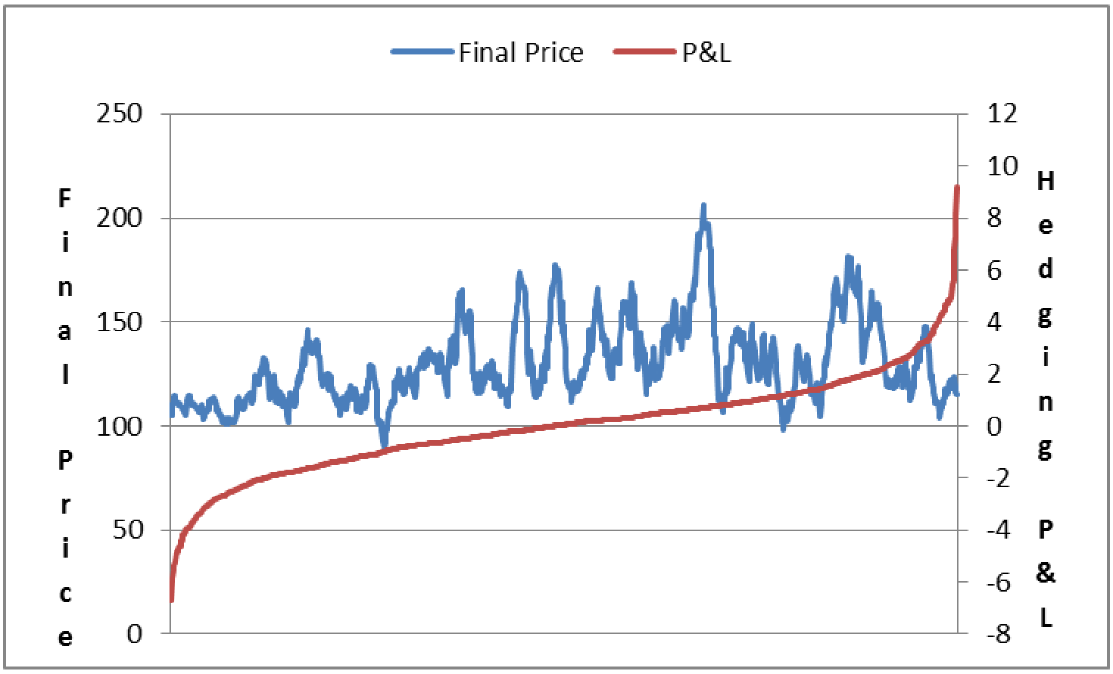

Note: this graph has been constructed by ordering the simulations in order of the final hedged P&L outcome and then plotting the final stock price against that statistic. A moving average across the simulations as so ordered was used to eliminate excessive noise.

In the middle of the P&L outcome range there is no correlation with the final price. But it can be seen that at the extreme negatives of hedging P&L that the final price is concentrated around the at the money price of 100[8]. This is due to the well-recognised fact that the second order greek, gamma, is at its highest when we are both at the money and close to maturity (see Appendix 1). A similar dampening of the spread of final price can be seen at the extreme positives of P&L though this is less pronounced because of the skewness in the distribution of realised volatility.

Note that it is not that finishing at the money implies that there will be realised volatility losses as there is an equal chance of profits. However as gamma is high the variability of the outcome is aggravated. This emphasises the importance of "maturity diversification" in stabilising hedging P&L in the overall context of a VA writer's portfolio.

[8] It rarely actually touches 100 because of the moving average technique I used to make the figure more readable

Underlying Volatility risk

The other obvious weakness in the hedging strategy is how do we determine the key parameter viz. the underlying volatility? In practice, we only have historic data to make estimates for this parameter but we must expect that this can only be a broad indication of what is correct for the future.

Figure 3 shows the sensitivity of the effectiveness of the hedging outcome to the underlying volatility being different from the assumed volatility in the hedging model.

Figure 3: Effectiveness of hedging for different underlying volatility[9]

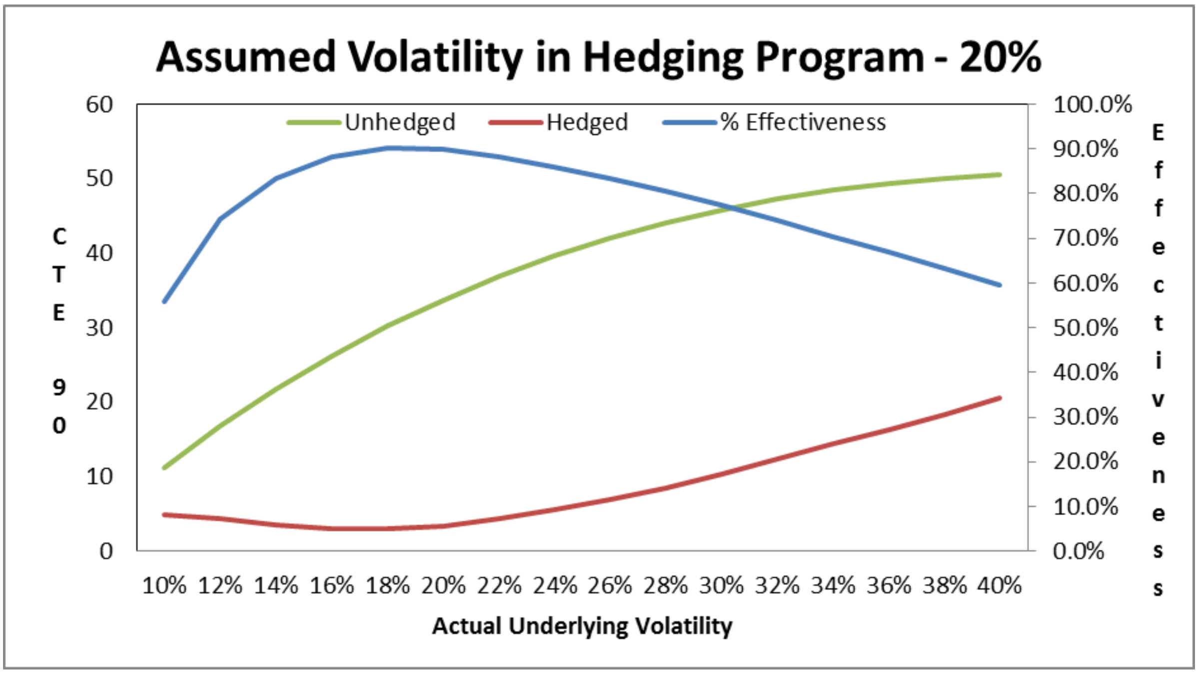

Note: the hedging strategy in the above figure uses a risk neutral 20% volatility. Hedge effectiveness is measured as the percentage reduction in the unhedged CTE90 brought about by the hedge.

As is to be expected, this strategy which is based on a 20% volatility assumption is at its most effective when the underlying volatility is around 20%.

The surprising aspect of this graph is that the green and red lines are broadly parallel when the underlying volatility is greater than the hedged volatility. This means that the absolute relief garnered from the hedging program, being the difference between the green line and the red line is quite insensitive to the "correctness" of our choice of hedge volatility.

The shape of the graph of the percentage effectiveness is a fairly symmetric hump shape as we might expect.

[9] In constructing this figure the underling simulated prices reflect the change in the assumed underlying volatility but use the same 60,000 uniformly distributed random numbers as are used throughout the paper.

Figure 4 below looks more closely at the situation where underlying volatility is 40%.

Figure 4: Comparing hedging programs with different assumed volatility

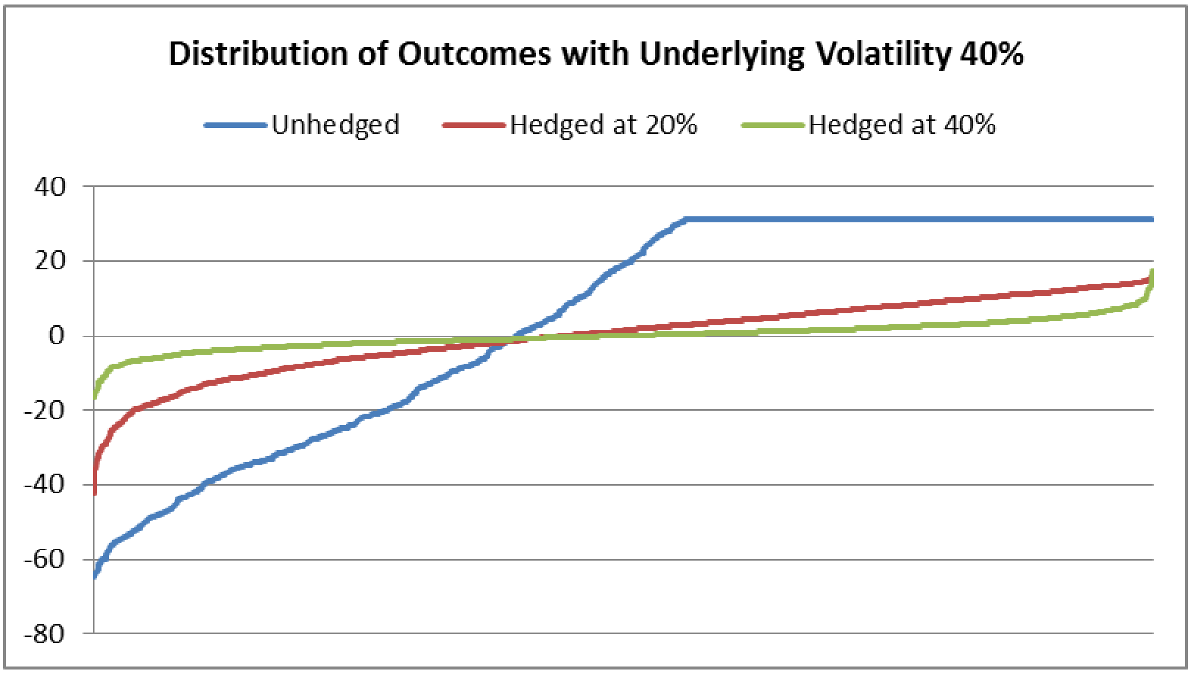

Note: Each line represents the 1,000 simulated outcomes ordered according to the particular P&L outcome which it represents. They do not therefore relate to the same ordering of the simulations. The initial value for determining P&L is taken in each case as the Black Scholes value using 40% volatility.

The 40% underlying volatility simulated in this figure is at the level of the CBI stress tests albeit over the full 5 year term and not simply years 1,3 and 5 of much longer liabilities. If we have hedged 20% volatility then it can certainly be said that we have got it fairly badly wrong and yet the hedging is still quite effective and not too far off the effectiveness of a hedging program which had the correct 40% volatility assumption, at least in a broad middle range of outcomes.

This is not saying that the final outcome is not sensitive to the underlying volatility; indeed the initial value of the Put option using 40% volatility is 28.3 compared to 12.5 value using 20% volatility. What it is saying is that the effectiveness of hedging in **reducing the dispersion** of final outcomes is relatively insensitive to having chosen the "correct" volatility assumption.

Put more colloquially, the good news is that it is unlikely that you would be seriously wrong in your choice of hedge formula but the bad news is that if the option is fundamentally mispriced there is no way to hedge your way out of that situation.

Table 3 below is a stark illustration of this phenomenon. In this Table I have assumed that underlying volatility is 40% whilst the IFRS reserves, the CTE gross capital and the hedge effectiveness have all been calculated using a volatility assumption of 20%.

Table 3: Allowance for hedging in capital calculation

	Correct Position	Actual Position
IFRS Reserves	28.3	12.5
Extra CTE90 unhedged capital	45.8	30.6
Allowance for hedging	-27.2	-27.5
Combined net reserves and capital	**46.9**	**15.1**

Note: The allowance for hedging in the "correct" position assumes we actually use 20% volatility in the hedge formula even though the correct assumption would have been 40%.

What the above Table demonstrates is that getting this key assumption wrong is highly significant in terms of calculating IFRS reserves and gross capital but the allowance for hedging effectiveness is not compounding the error in any significant way. I have used a hedge effectiveness factor of 90% in calculating the net capital but I have applied it to an unhedged capital calculated at 20% volatility. Both assumptions are wrong but self-cancel to render the absolute deduction for hedge effectiveness more or less correct. It should be noted that in practice calculate their capital requirement using a prudent estimate of underlying volatility, for example, 80th centile of historic realised volatility.

Targeting Volatility

The considerations in the last section lead to the concept of targeting volatility. A typical algorithm would calculate the annualised realised volatility over, say, the last 20 days and adjust the exposure to the stock or index so that if this were indeed the underlying volatility going forward the volatility of the adjusted exposure would be the target volatility. There would usually be a "cap" applied to the calculated exposure. A typical algorithm would therefore be as follows:

$$Exposure\ to\ stock/index = Min\left({Target\ Volatility}/{Realised\ Volatility_{20\ days}}, Cap \right)$$

From the observations in the last section it should be noted that this mechanism has a relatively modest impact on hedging effectiveness. Its main purpose is to minimise underlying volatility risk in the pricing itself.

The rationale for a target algorithm could either be that there is too much uncertainty in the future underlying volatility or else that we expect it to vary with time. The assumption is that recent realised volatility will be a likelihood estimator of underlying volatility in the immediate future. In Table 4[10] we are targeting 20% volatility using a 20 day algorithm.

Table 4: Effect of a target algorithm in adjusting underlying volatility

Cap	Underlying Volatility		
	15.0%	**20.0%**	**30.0%**
100.0%	*15.0%*	*19.1%*	*20.9%*
110.0%	*16.4%*	*20.0%*	*21.0%*
120.0%	*17.7%*	*20.5%*	*21.0%*
130.0%	*18.8%*	*20.7%*	*21.0%*
140.0%	*19.6%*	*20.9%*	*21.0%*
150.0%	*20.1%*	*20.9%*	*21.0%*
No Cap	*21.0%*	*21.0%*	*21.0%*

Note that if there is no cap in place the adjusted underlying volatility is slightly greater than the target volatility of 20%. That is because the algorithm itself contributes an element of volatility.

The cap has minimal effect when underlying volatility is high or even equal to the target volatility. However, if the underlying volatility is low the cap can have a fairly significant effect in undershooting the target volatility.

[10] The table was constructed using Monte Carlo methods as the presence of the cap makes analytical methods somewhat intractable

What to hedge

At this stage I take a diversion and look at the so called "Stop/Loss" hedging strategy[11]. Under this strategy the delta is calculated by reference to the intrinsic value of the option i.e. the pay-off if the stock price at the maturity of the option was the current price. The intrinsic value by definition satisfies the boundary conditions so if we can successfully delta hedge it we should be alright. When out the money the intrinsic value is zero and so too is the delta. When in the money the delta is negative unity as each euro fall in the stock increases the intrinsic value by a euro as can be seen in Figure 5.

Figure 5: Stop Loss strategy for a Put option

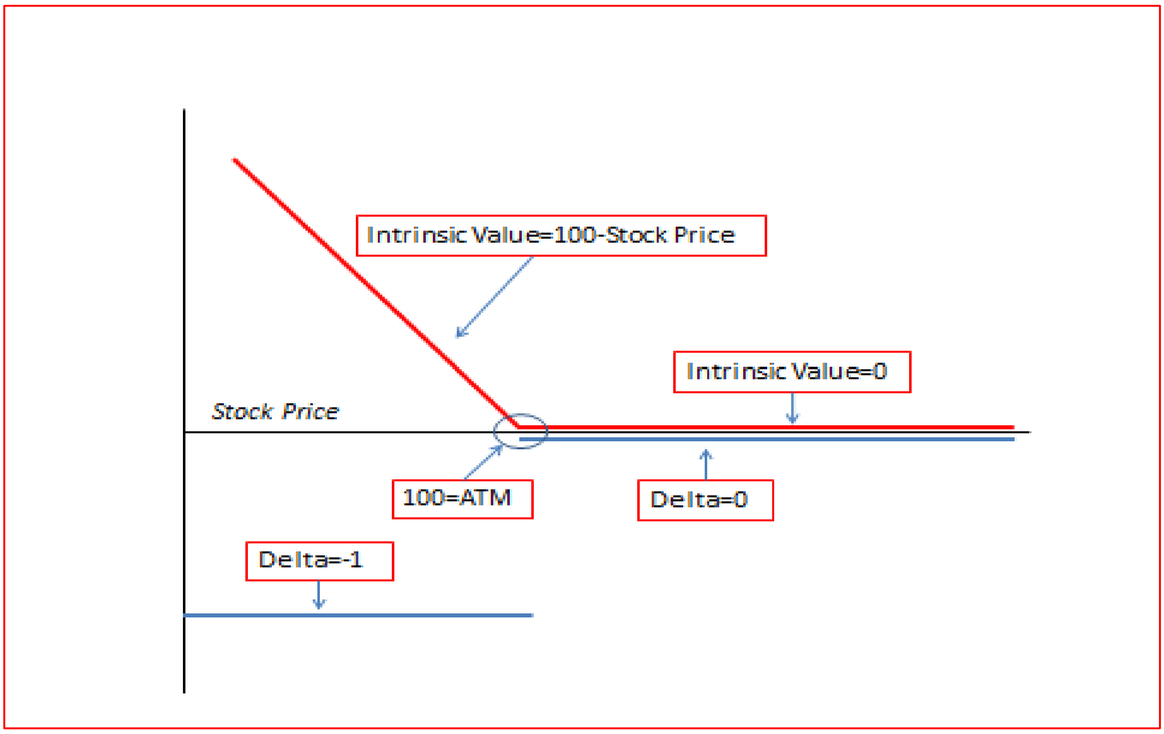

This is equivalent to using a delta hedging strategy with zero[12] volatility. The second derivative i.e. the convexity (aka gamma) would appear to be everywhere zero and since we know that convexity losses are the main source of variability in a delta hedging strategy we may be surprised that this stop-loss strategy is almost as inefficient as an unhedged strategy, producing a CTE90 of 27 or 82% of that for the unhedged CTE90. The conundrum is explained by noting that all the convexity is concentrated at the "at the money" point. The second derivative is infinitely positive at this point. The convexity losses are picked up in those simulations which oscillate a lot about the at the money point.

The above strategy is therefore easily dismissed as a practical proposition but there are a few alternatives that could plausibly be used instead of the IFRS Risk Neutral strategy discussed above. An obvious alternative candidate would be to hedge the Real World option value[13].

[11] This section draws on similar comments in Hull pp 310-311
[12] An unhedged strategy is equivalent to hedging infinite volatility.
[13] Meaning here using the real world equity return in the Black Scholes option formulae. Strictly speaking a Real World option value is not a valid concept as consideration of the discount rate forces one towards construction of a replicating portfolio and thus toward the risk neutral martingale framework.

Yet another candidate would be to hedge the CTE90 of final outcomes. These three dynamic future hedging strategies can be categorised broadly as follows:

Table 2: Description of Hedging Strategies

Strategy	**Broad description**
Risk Neutral Option Value	The market consistent value of the option and what is used in IFRS accounts
Real World Option Value	The actual expected value of future outcomes albeit discounting using the Real World expected return on the stock
CTE90	The Solvency I capital requirement

For completeness I also look at the results for a static delta hedge i.e. maintaining the opening short position on the hedging program which is 33% of the nominal value of stock in the Put option; this would involve continually rolling over the short futures.

Figure 6: Overall P&L outcomes for various hedging strategies

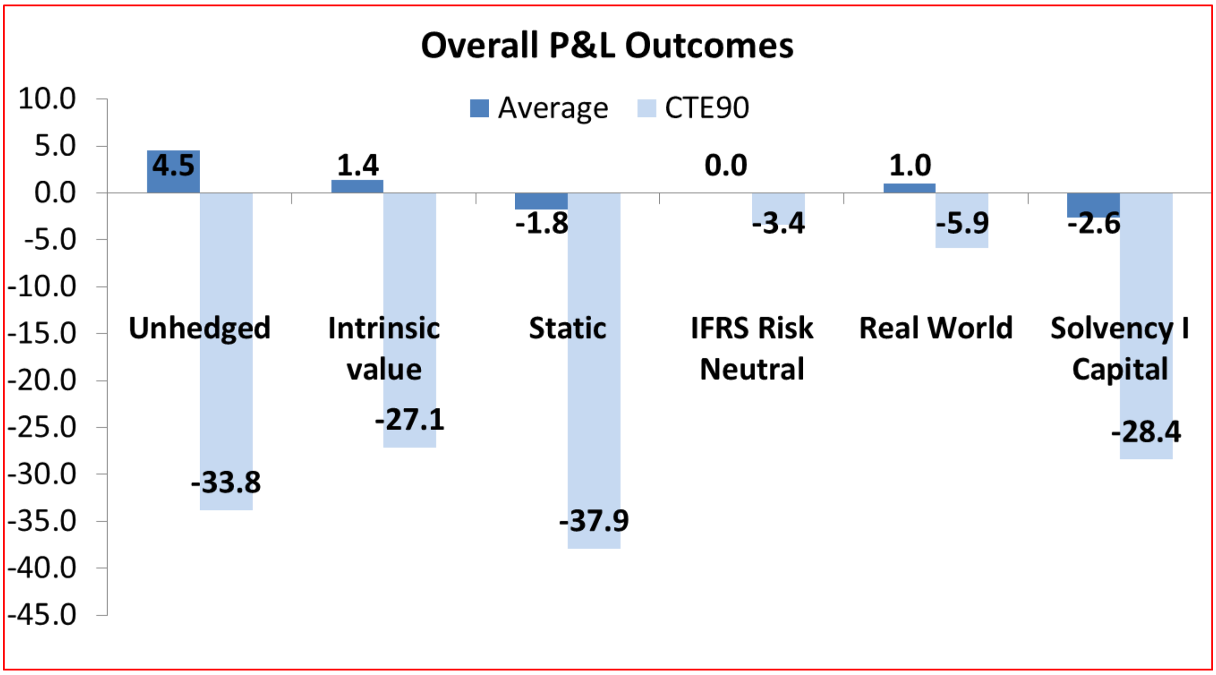

Figure 6 shows how these various strategies fare.

The average final outcome is relatively insensitive to choice of hedge strategy but reflects the extent to which a particular strategy is shorting the equity risk premium. Since the opening position in each case is the Risk Neutral option value the Risk Neutral strategy is neutral.

The more interesting statistic is the CTE90 worst outcomes. The two expected value hedging strategies are dramatically effective in reducing this CTE90 with the Real World strategy nearly as effective as the Risk Neutral strategy but not quite.

The Solvency I Capital hedging strategy gives similarly poor results to the Stop/Loss or Intrinsic Value strategy.

The static strategy is even worse than the unhedged strategy. It is noteworthy that the static strategy does dampen the 10% worst unhedged scenarios producing a conditional expectation of 14.7 for those scenarios, thus apparently being almost 60% effective in reducing the unhedged CTE90 of 33.8. However, whist the static hedge is this effective in these particular scenarios it is others which dominate the worst 10% of its own scenarios. The message here is that an effective hedging strategy has to be dynamic.[14]

[14] I ignore the availability of derivative capital market instruments which might more closely match the liability as ultimately these need to be delta hedged in the primary market. Nonetheless, especially when it comes to interest rate optionality there can be a role for swaptions in dampening the need to dynamically rebalance.

The One Year View

Up till now I have considered the variability of the final outcome. The final outcome, once it has been reached is of course an absolute. The outcome over one year requires some measure at the end of the year of the potential final outcome. In IFRS accounting this is usually the market consistent Risk Neutral value of the option. Figure 7 shows the effect on the IFRS outcome over one year of various hedging strategies.

Figure 7: One Year IFRS P&L of different hedging strategies

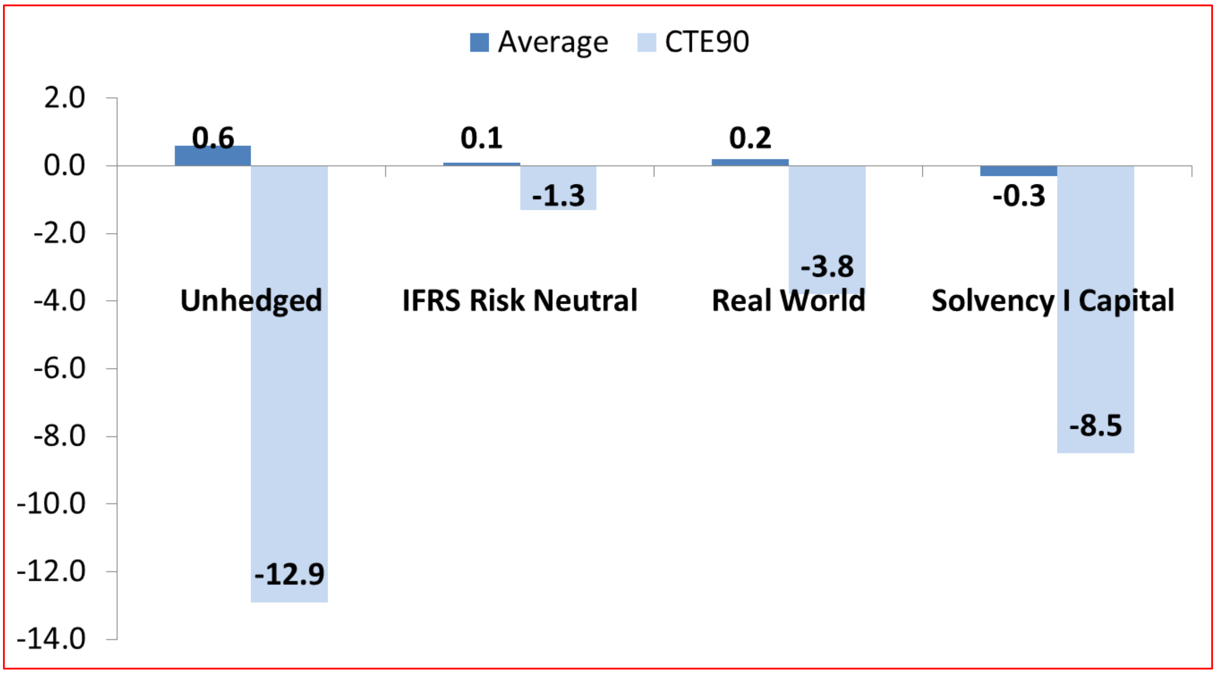

The picture is directionally very similar to the same figure for the final outcomes albeit the figures are scaled back. The gap between the Real World CTE90 and the Risk Neutral CTE90 has widened somewhat which is not surprising since we are actually measuring the outcome using the Risk Neutral value of the option.

Another measure which we might be interested in would be the Solvency I Capital. This is the CTE90 requirement set by the CBI for VA business sold before 2011. The picture for this one year measure is shown in Figure 8.

Figure 8: One Year Solvency I capital changes under different hedging strategies

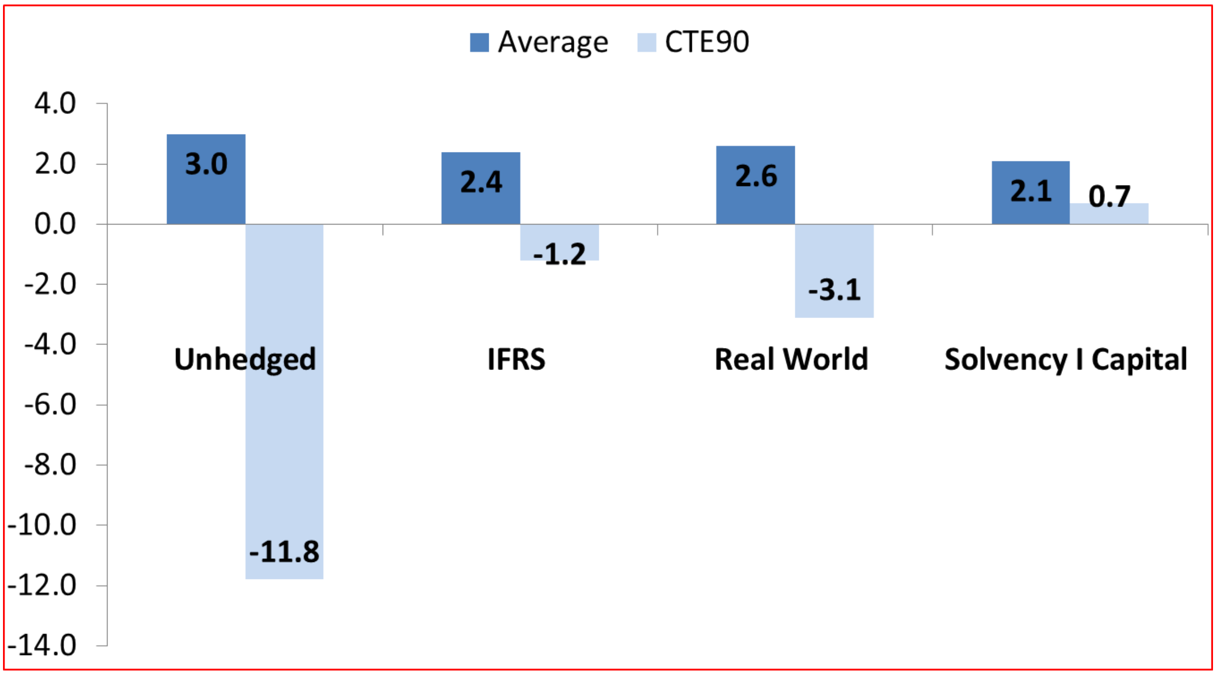

The average position in all cases is a fairly significant positive outcome with again the relativities being based on the comparative "shortness" of the strategy. This reflects the release of solvency capital as one gets a year closer to the maturity of the option.

Not surprisingly the Solvency I Capital hedging strategy is the most effective at dampening the dispersion in its own outcome with the one year worst CTE90[15] of this strategy being an actual release of 0.7 of capital.

[15] Just to be clear the measure which is being examined over one year is the change in the Solvency I capital requirement which is itself a CTE90 of final outcomes. The table shows the one year worst CTE90 of this measure.

Historic Effectiveness of Hedging Strategy

The effectiveness of a hedging strategy is not known until the final outcome. We can, using our model, simulate what the distribution of this measure might be. Hedging Effectiveness is defined by the following formula:

$$Historic\ Hedging\ Effectiveness = \left[1 - \left|\frac{Hedged\ P\&L}{Unhedged\ P\&L}\right|\right]^{+} \qquad \text{(A)}$$

Note that this is a somewhat different concept from that used in Figure 3. Figure 3 is a prospective statistical measure; it compares the CTE90 of the hedged strategy with the CTE90 of an unhedged strategy and is the concept most relevant in setting capital requirements. The historic hedge effectiveness is a measure of how much the P&L has been dampened by the effect of the hedging strategy; it should be measured as the percentage change in the absolute levels of the respective P&Ls. A 100% effective hedging strategy would produce zero P&L.

Figure 9: Simulated historic hedge effectiveness over the full term

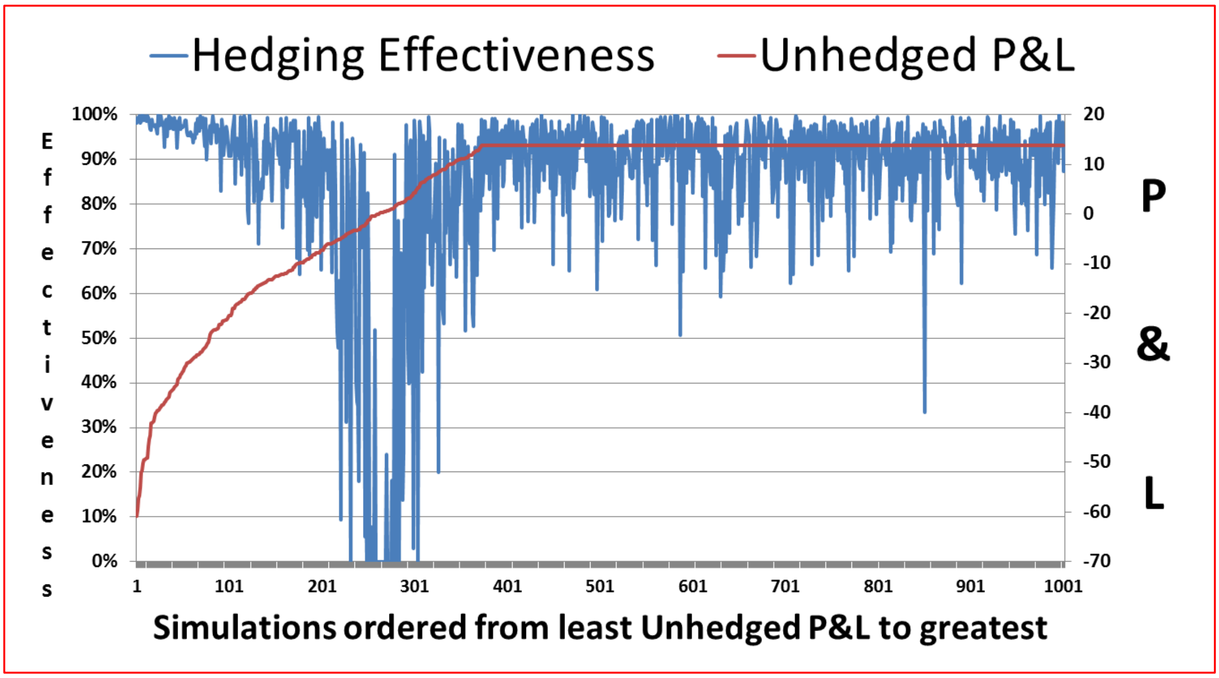

Figure 9 shows the simulated historic hedge effectiveness over the full term of the option assuming underlying volatility and hedge volatility were both 20%. The simulations have been ordered from least unhedged P&L outcome to greatest and we note that, as expected, it finishes as a horizontal line at a height equal to the accumulated initial value above the X axis. The hedge effectiveness as calculated by the above formula has been plotted against these simulations.

From Table 2 we know that the CTE90 hedge effectiveness is around 90%. From Figure 9 we see that we will only be reasonably assured of observing this level of historic hedge effectiveness if the outcome has in fact finished in the “tails”. If the final outcome is somewhere in the middle the observed historic hedge effectiveness can be anywhere. Indeed we note that if the actual unhedged outcome is near breakeven we will get negative historic hedge effectiveness and since this has been zeroised we see a little gap in the figure.

As there is little or no experience of the final outcomes of actual VA hedging strategies the term "historic effectiveness" is invariably reserved for measuring the historic effect in stabilising the change in a prospective measure, usually the IFRS LOV but also possibly the capital requirement. The following is a quote from the CBI's recent paper on stress testing VAs.

"The effectiveness of the hedging program at offsetting total guarantee liability movement in the base scenario is not to exceed that based on a credible amount (2 years or more) of historical data."

I presume that the CBI is indicating that it measures effectiveness by the historic ability of the hedging program to neutralise changes in LOV caused by the realised changes in the underlying stochastic drivers. In our case the only stochastic driver is the stock price. Figure 10 shows the effectiveness of the hedging strategy on dampening first year IFRS LOV outcomes in each of the 1,000 simulated scenarios.

Figure 10: One year hedge effectiveness plotted against unhedged P&L

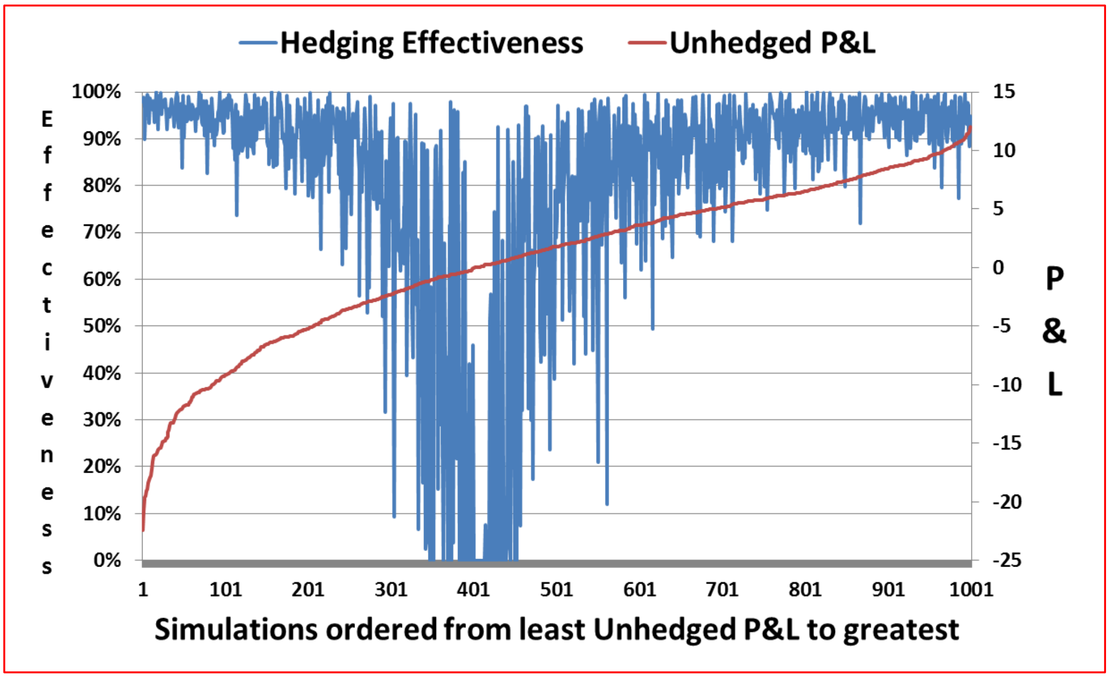

Note: the underlying volatility is taken as the base case 20% and the hedging strategy is Risk Neutral also with 20% volatility

Figure 10 is similar to Figure 9 in that it tells us is that this measure of effectiveness is really only credible at the extremes of unhedged P&L. The average historic hedge effectiveness on the one year view is 80% and its standard deviation is 26%. It is difficult to see how two observations could give anything which could be described as a credible indication of the "correct" result which we know from Table 2 is 90%.

In Figure 11 below, using the same ordering of the simulations, I show the probability of observing a hedge effectiveness of over 85% in one year's P&L outcome.

Figure 11: Probability of observing one year hedge effectiveness greater than 85%

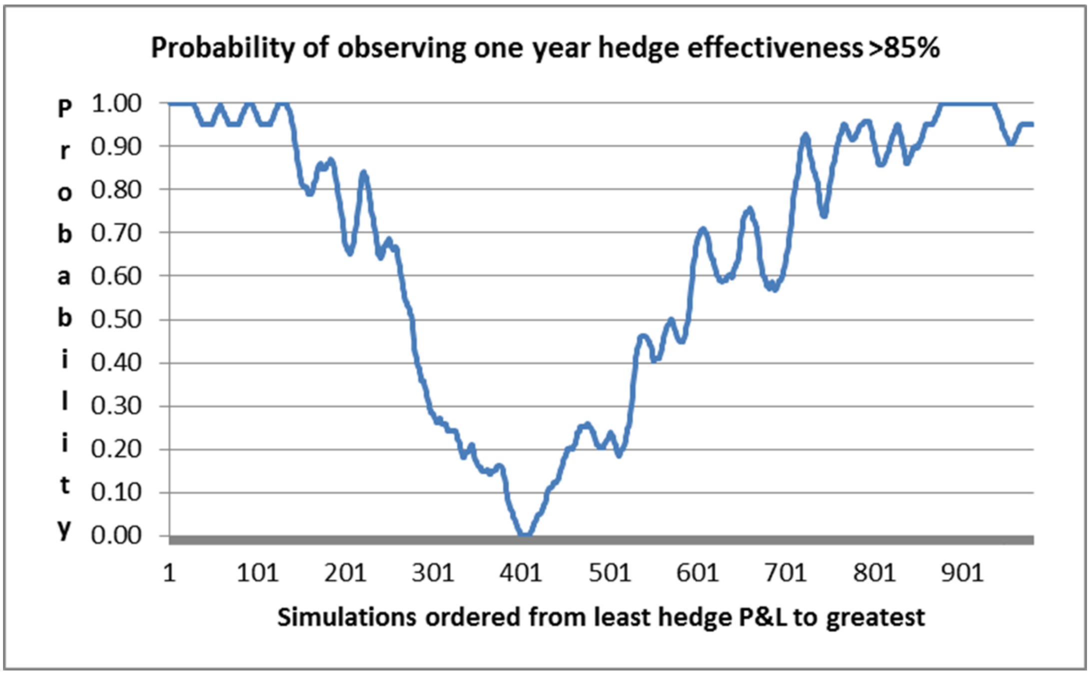

Note: Underlying volatility and hedging volatility are both 20%. Excessive noise has been removed by the technique of moving averages across the simulations as ordered.

In recent times experience has been somewhat in one extreme and therefore actual company hedging P&Ls are showing plausible values for effectiveness. The situation is also considerably complicated in that durational interest rate matching is usually covered within the same hedging paradigm using the so called “rho” sensitivity to determine appropriate swap hedges. Since a good swap hedge will also inherently include an interest rate convexity hedge this element of the hedging strategy will always appear highly effective. This further complicates inferring the “optionality” hedge effectiveness from historic results albeit it biases the observation towards overstatement whilst the random nature in the above figures would tend to bias it towards understatement.

It is not unreasonable for the CBI to want the allowance for hedge effectiveness in the CTE calculation to be underpinned by historic metrics from the hedging program. The above demonstrates that the crude hedge effectiveness factor in equation (A) is not fit for that purpose.

P&L Attribution

Let us recall the key Taylor expansion.

$$\delta Lov = \Delta\delta s + \Theta\delta t + \frac{\Gamma}{2}(\delta s)^2 + o(\delta t) \qquad (1)$$

Taking expectations of this equation and using the bra-ket (<>) convention for expectations we get:

$$< \delta Lov >= \Delta < \delta s > + \Theta\delta t + \frac{\Gamma}{2} < (\delta s)^2 > + o(\delta t) \qquad (2)$$

Subtracting (2) from (1) we get:

$$\delta Lov - < \delta Lov >= \Delta(\delta s - < \delta s >) + \frac{\Gamma}{2}((\delta s)^2 - < (\delta s)^2 >) + o(\delta t) \qquad (3)$$

Where we note that the Θ terms cancel out but we still have a $o(\delta t)$ term as this residual term is not necessarily the same quantity in (1) and (2).

If we note that the expectation of overall P&L is zero then analysing the difference between actuals and expectations of the various constituents of P&L will give an attribution of it.

Thus from (3) and assuming a full delta hedge is in place so that at the net balance sheet level $\Delta = 0$ we get the key P&L attribution formula:

$$P\&L = \frac{\Gamma}{2}((\delta s)^2 - < (\delta s)^2 >) + o(\delta t)$$

Or in words:

$$P\&L = \frac{Gamma}{2}(Realised\ Volatility - Expected\ Volatility) + Residual\ Term$$

The residual "unexplained" term tends to cause great angst as we do not like to see things unexplained. In this simple case it would of course have been "explained" if we had used the higher order terms of the Taylor expansion as there is nothing else at play. Figure 12 illustrates the situation.

Figure 12: Unexplained/residual hedge P&L

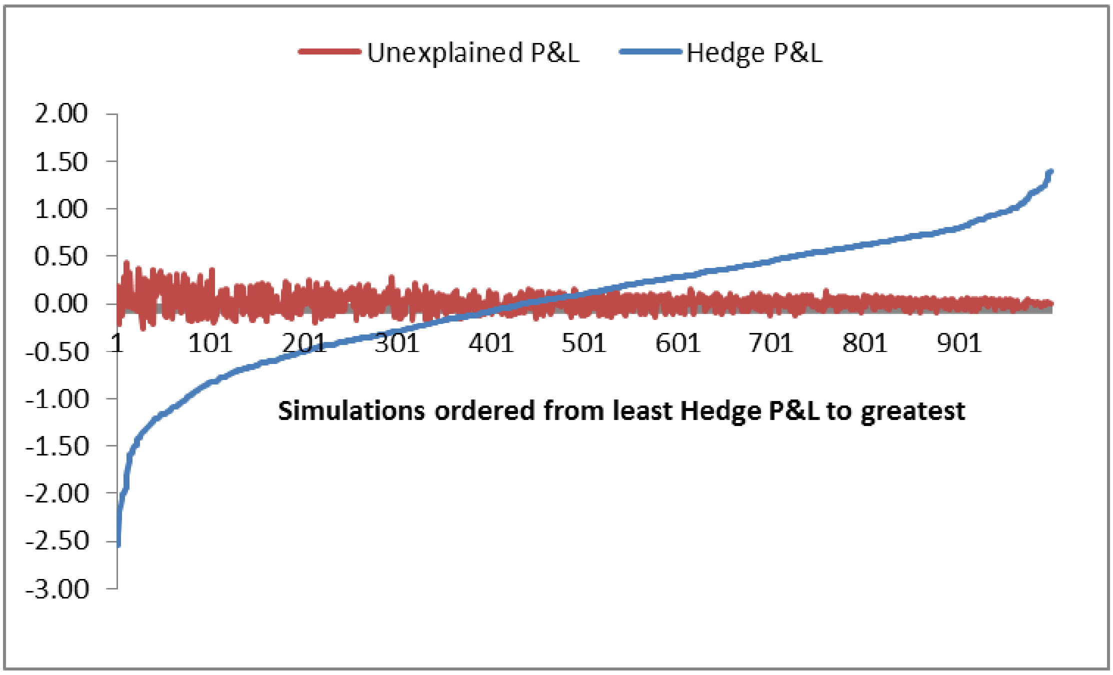

Again here the technique is to order the simulations by Hedge P&L from lowest to highest. Note how the "unexplained" tends to be greatest in the loss making situations which is not surprising as these are where the explained losses are themselves highest due to high realised volatility. We see that in this very simple situation the Taylor higher order noise or the unexplained by gamma alone is not insignificant being of the order of 25bp of the nominal amount of the host contract. And of course when P&L itself is close to zero we get the irony that the unexplained appears very high in relation to that P&L.

We can expect that the Taylor higher order effect would reduce in relation to the gamma effect with increased frequency of rebalancing using the usual square root adjustment and since in practice re-balancings are much more frequent than monthly it is possible that in practical VA situations the higher order Taylor effects are indeed negligible.

The CBI Stress Tests

Figure 13 illustrates the effect of the CBI stress tests on this simple contract at its inception.

Figure 13: The CBI stress tests

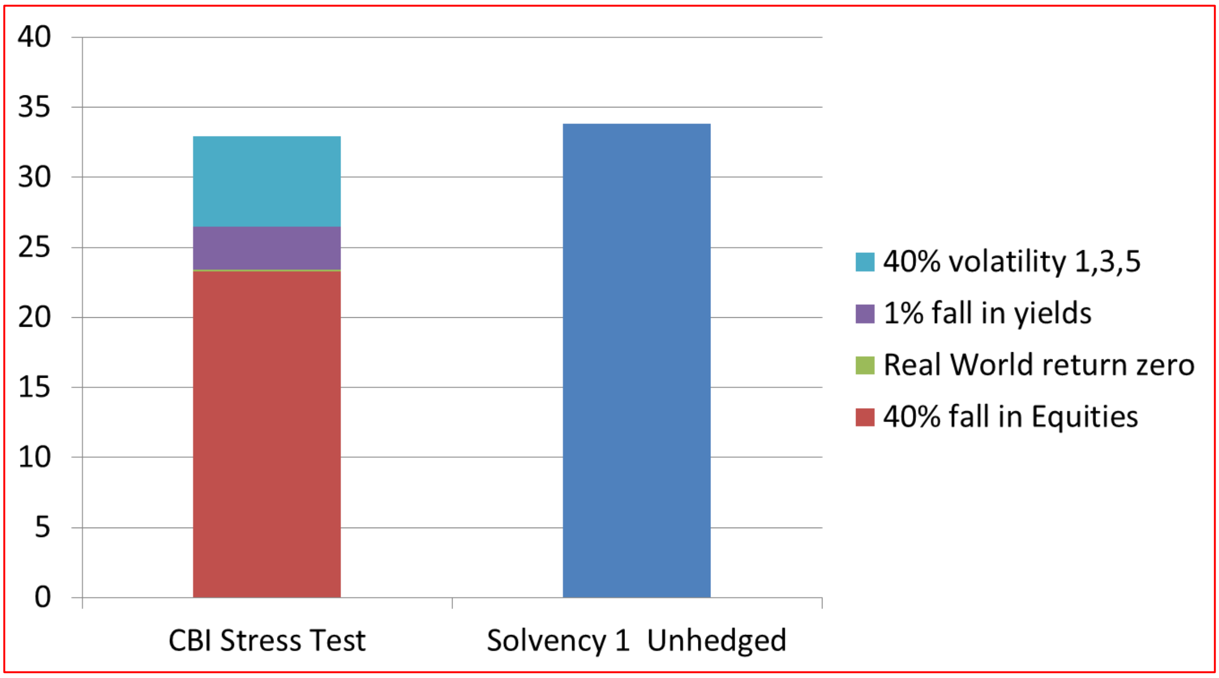

The cumulative effect is very similar to the unhedged CTE90. They effectively eliminate any credit for hedge effectiveness, at least for the simple example in this paper.

I have interpreted the stress test for volatility as referring to the *underlying* volatility and the above figure shows the average of the simulated outcomes, given these stressed parameters.

An alternative approach would be to assume that the stress test refers in a quasi-deterministic way to *realised* volatility. Further assumptions are then needed such as that the price changes in a deterministic up/down fashion with uniform jumps at rebalancings or, alternatively, gamma could be estimated from other aspects of the model. I expect the results to be similar to the stochastic approach which I took above.

In passing, I allude to the G1 stress test which reads as follows:

"10% proportionate reduction in the effectiveness of the hedging programme at offsetting the total guarantee liability movement for the full projection term in addition to the solvency II capital charge for operational risk at 31/12/2012"

The stress appears to be applied to the economic balance sheet and as this will not actually reflect the effects of the future hedging program I found it difficult to interpret for my example. In a Solvency I context, which does allow for the effects of the hedging strategy, this is more meaningful and would amount to a stress of c. 3.0 in the example.

Solvency II

Figure 14 illustrates my understanding of how the current Standard Formula would apply in this simple example.

Figure 14: Solvency II standard formula

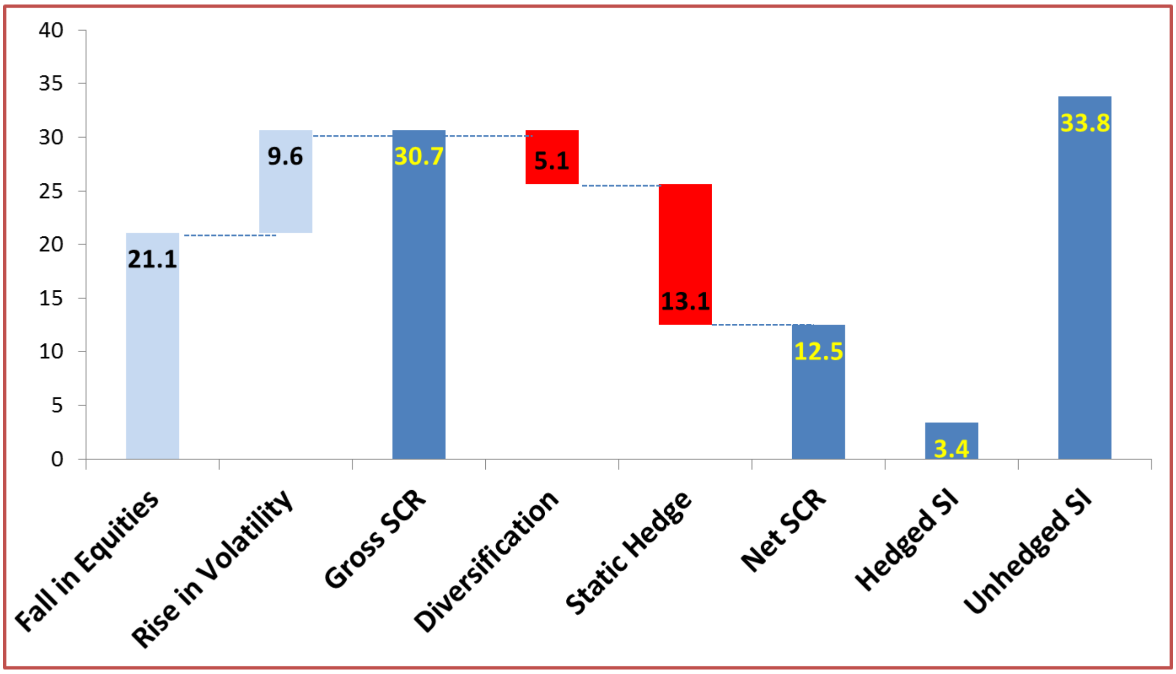

The gross SCR i.e. ignoring the hedge instruments and diversification, is close enough to the Solvency I unhedged CTE90. The hedged Solvency I is in this example exaggeratedly low because of the 90% hedge effectiveness. The Net SCR is equivalent to 63% hedge effectiveness which is at the lower range of that shown in Figure 3.

Noteworthy is the significant relief from the current hedging position. This is at variance with the evidence of Figure 6 which actually indicated that the current balance sheet hedge position, if it is maintained, is at even higher risk than a fully unhedged position. Of course, given the complexity of VA products most writers will not be using the Standard Formula.

Practical Afterthoughts

Throughout this paper I have been able to indulge in the simplifications of the Black Scholes ivory tower. In practice many complications present themselves of which the following list is far from exhaustive.

- VA products invariably contain demographic and behavioural dimensions which, as well as being unhedgable, considerably complicate the hedging models.
- Generally VA products are much longer term than five years and may even allow regular premiums or future top ups.
- Given the complications, closed form solutions would rarely be utilised. Instead simulations are performed, perhaps nightly, not only to evaluate the LOV but more importantly its attendant "greeks" which feed into the hedge trading platform. A greek is calculated by running the simulations again for each parameter being hedged by slightly tweaking that parameter in isolation and using the difference between that run and the base run to approximate the greek.
- Interest rate played a small part in the example in the paper but for much longer term products it becomes more important and in the case of a GMIB[16] it becomes especially significant as it now includes optionality. In particular, there is not just one interest rate but in practice hedging programs might split up the yield curve into say 10 durational "buckets".
- Basis risk can be present whereby the hedging tradable does not precisely match the fund to which the options apply.
- Volatility targeting has been mentioned in the paper and presents its own complications.
- Higher order greeks such as gamma can be hedged but it should be noted that there are no natural buyers of convexity risk in the way that there are natural buyers and sellers of the primary instruments. Selling convexity risk is really just passing the parcel.
- P&L attribution can in practice be particularly intricate with such things as the second order "cross greeks" playing a role.
- Product charges must be allowed for and have themselves a stochastic dimension.
- Last but not least dynamic hedging incurs transaction costs.

[16] GMIBs guarantee a level of life income and are akin to Guaranteed Annuity Options.

Appendix 1: Illustration of Convexity Effects

I find the following visualisation a useful demonstration of why we always have convexity increases in an option value.

Figure 15: Illustration of the cause of convexity of option values

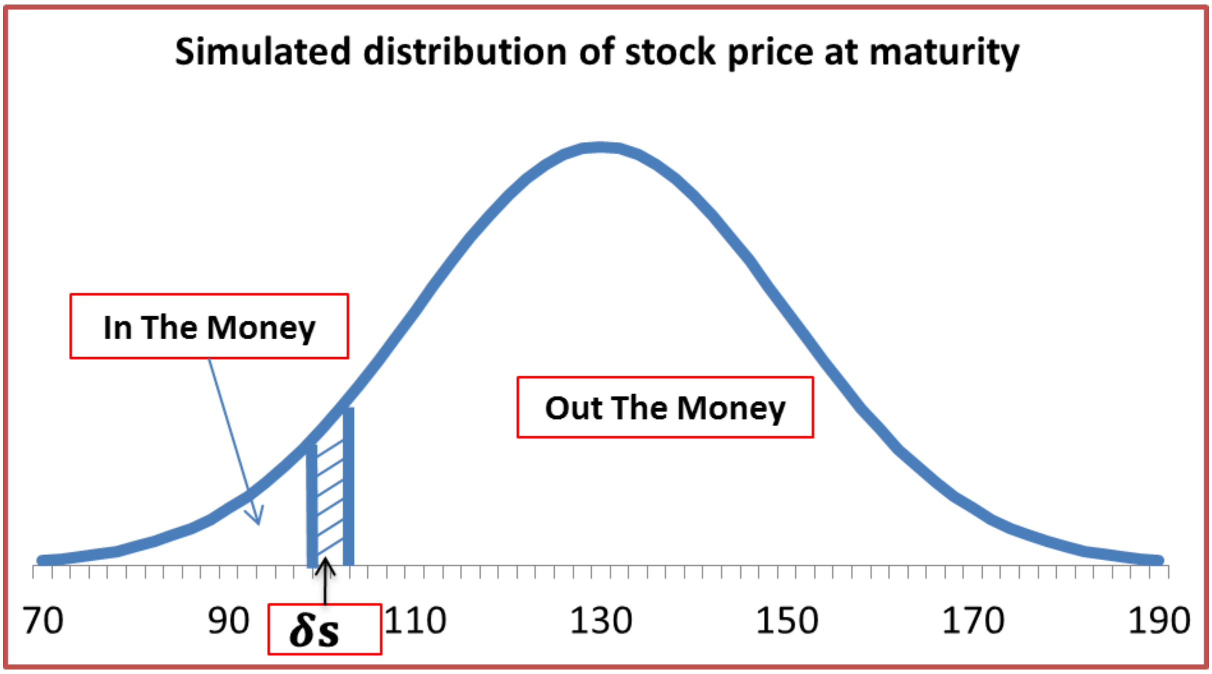

The bell curve[17] in the diagram represents the simulated distribution of the final maturity price at some point in time prior to maturity.

The At The Money Point is 100. Those simulations which are to the left of this point are "in the money" and those simulations to the right are "out the money". The value of the option under a risk neutral measure for the bell curve is (ignoring discount for simplicity) the sum of the intrinsic values of each simulation. The intrinsic value to the left of the ATM point is "100 – S" whilst to the right it is 0. We can see that the delta of the aggregate of these value is the (negative) cumulative distribution function up to the ATM point. Gamma, being the derivative of delta, is the probability density function at the ATM point[18].

A delta hedge effectively follows a stop-loss strategy for each simulation. Simulations to the left of ATM are fully hedged and to the right have no hedge in place. To more clearly follow the rest of the explanation imagine that such a delta hedge is in place.

[17] I have ignored the "log" nature of the actual price model and produced a symmetric bell curve. It does not impact on the illustration.

[18] Whilst delta is negative a double negative makes gamma positive. For a Call option both delta and gamma are positive, delta being the sum of "S – 100" for the simulations to the right of the ATM point and gamma being the same as for the Put option.

Now, if we have a small change δs in the stock price the bell curve will make a small shift either to the left or the right. If it is a price fall the curve shifts to the left and that means that the simulations within the narrow strip in the diagram have moved from OTM to ITM. As they have not been hedged the writer of the option loses on average $\frac{1}{2}\delta s$ on each of these simulations. The area of the strip represents the number of simulations involved and its size is $\Gamma\delta s$ so that gives a total loss of $\frac{\Gamma}{2}\delta s^2$.

Similarly if the small change was positive the curve will shift to the right and the strip now represents simulations which we have hedged but which now enjoy the Put floor. In other words we continue to lose on the hedge instruments backing these simulations without a corresponding fall in the liability.

We can also see that the closer we are to the ATM point the greater is gamma. Also as we approach maturity, the bell curve tightens (we are more certain of the final outcome) with it becoming more peaked. So the most dangerous combination for the hedging strategy is to be at the money approaching maturity. That will not necessarily mean that we will make losses. To make losses we need the realised volatility to exceed the expected volatility, but it does mean that the variability of the outcome has increased.

Appendix 2: Black Scholes and other formulae

Black Scholes formulae for a Put option[19]:

Value of Put option (p)

$$p = Xe^{-rt}N(-d2) - SN(-d1)$$

where:

X is exercise price

r is the risk free rate

t is time to maturity

$$d2 = \frac{\ln\left(\frac{S}{X}\right) + \left(r - \sigma^2/_2\right)t}{\sigma\sqrt{t}}$$

$$d1 = d2 + \sigma\sqrt{t}$$

and *N()* is the cumulative standard Normal distribution. Note that *N(-d2)* is the probability of the option being in the money at maturity.

Delta of Put option (Δ):

$$\Delta = -N(-d1)$$

This might seem rather obvious from the above formula for *p* but it should be noted that *d1* and *d2* are both dependent on *S* and it is somewhat coincidental that the Δ can be calculated as if they did not have that dependence

Gamma of Put option (Γ):

$$\Gamma = \frac{N'(d1)}{S\sigma\sqrt{t}}$$

where *N'()* is the derivative of the cumulative standard Normal distribution i.e. the probability density function of that distribution as follows:

$$N'(x) = \frac{1}{\sqrt{2\pi}}e^{-\frac{1}{2}x^2}$$

Closed Form Formulae for the CTE90 of a Put option

Value:

The present value of the CTE90 of a Put option is given by the following integral:

$$10 \times e^{-rt}\int_{-\infty}^{S_{.1}^M} p(s)(X-s)^+ dx$$

where *p(s)* is the p.d.f. of the value of the stock, s, at maturity, and *r* is the discount rate

$S_{.1}^M$ is the 10th centile of the maturity value, and equals:

[19] Hull pp241 – 242, 326

$S\left(\exp\left(r - {\sigma^2}/{2}\right)t - \ z_{.9} * \sigma\sqrt{t}\right)$, where $z_{.9}$ is the 90th centile of the standard Normal distribution

$(X - s)^+$ is *X* - *s* if *X* > *s* otherwise 0

The factor of 10 reflects the conditionality as the integral of the p.d.f. over that range is, by definition, 0.1.

The situation is complicated by whether one is evaluating a Risk Neutral market consistent value or a Real World value which is more associated with CTE90. For the purpose of this paper I used Real World *r* both for the p.d.f. but also for the discount rate as this makes the math a bit more tractable.

To evaluate the above integral we recognise two possibilities:

Possibility 1: $S_{.1}^{M} > X$

In this case the discounted integral is simply the value of the Put with exercise price X

Possibility 2: $S_{.1}^{M} < X$

In this case the value of the discounted integral is the value of a Put with exercise price $S_{.1}^{M}$ plus $.1e^{-rt}$(*X* - $S_{.1}^{M}$) so that the value of the CTE90 becomes:

e^{-rt}(*X* - $S_{.1}^{M}$)+ 10 x value of Put option with exercise price $S_{.1}^{M}$

In summary, the value of CTE90 of a Put option is given by the following formula:

$$10 \ \times Put\left(exercise\ price = Min(X, S_{.1}^{M})\right) + e^{-rt}(X - S_{.1}^{M})^{+}$$

Delta:

Again we recognise two possibilities.

If $S_{.1}^{M} > X$ delta of the CTE90 is simply 10 times the delta of the Put with exercise price *X*.

If $S_{.1}^{M} < X$ we have two parts. The derivative of $(X - S_{.1}^{M})$ is minus the derivative of $S_{.1}^{M}$ with respect to *S* i.e. $-\exp(\left(r - {\sigma^2}/{2}\right)t - \ z_{.9} * \sigma\sqrt{t}))$.

To calculate the delta of the option part we revert to the formula for a Put option and we note that in this case both *N(d1)* and *N(d2)* are independent of *S* as *ln*(*S*/$S_{.1}^{M}$) is independent of *S*. However *X* is now dependent on *S*. The delta of the option part is therefore:

$$10 \times (\mathrm{e}^{(\mathrm{r}-.5*\sigma^2)\mathrm{t}-z_{.9}*\sigma\sqrt{t}} \times e^{-rt}N(-d2) - N(-d1))$$

combining the two parts we have the delta of the Put option when $S_{.1}^{M} < X$ as follows:

$$10 \times (\mathrm{e}^{(\mathrm{r}-.5*\sigma^2)\mathrm{t}-z_{.9}*\sigma\sqrt{t}} \times (e^{-rt}N(-d2) - .1e^{-rt}) - N(-d1))$$

Note that this formula has no dependence on *S*. However *S* does determine when we switch over from using this constant delta to using the one applicable when the 10th centile maturity price is greater than the exercise price.
Figure 16 illustrates the situation at inception.

Figure 16: CTE90 of Put option at inception and its Delta hedge

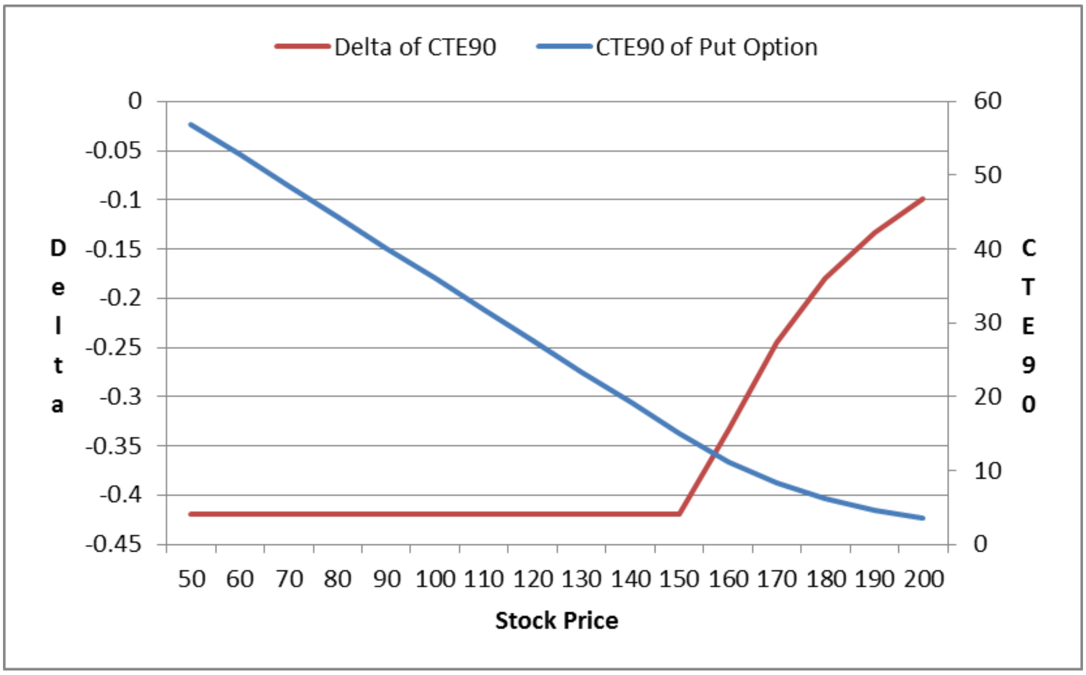

We can see that Figure 16 bears similarities to Figure 5 for the Stop/Loss strategy.

Appendix 3: A generalised approach to the attribution of stochastic P&L

The objects of our interest are Liabilities and Assets. I will generically denote a typical object by the letter $\mathcal{L}$ by analogy with the Lagrangian[20] of classical mechanics, although the object could equally be an asset or a liability.

We have a configuration space of q_is and t, where, to avoid getting too abstract, we can regard t as time. The q_is will typically be various index values or prices and various interest rates, yields or swap rates. In keeping with the analogy we will sometimes refer to them as co-ordinates.

Each q_i has a stochastic relationship with respect to t but of course t cannot be stochastic with respect to itself[21]. We are interested in how our objects evolve with t, so we see that t stands out in several ways from the other parameters of the configuration space.

Each object has associated with it a Lagrangian (for want of a better word) which determines its value and which has as many derivatives as we may need, as follows:

Property A: $\mathcal{L} = \mathcal{L}(q_{1,}q_{2,}q_3 \dots, q_{n,}t)$

In practice the LOV is the Lagrangian of most interest. It is invariably calculated from stochastic simulations of the liability outcomes based on an ESG model. Provided the number of simulations is large enough and the ESG is reasonably well behaved we can expect, by the law of large numbers, that the LOV will approximately have the required level of smoothness.

In general the q_is need not follow a geometric Brownian motion but we do require their variation to be of the order of $\sqrt{\delta t}$:

Property B: $\delta q_i = O(\sqrt{\delta t})$

It follows from these properties that the Taylor expansion of any Lagrangian to first order of smallness in δt is as follows:

$$\delta\mathcal{L} = \sum_i \frac{\partial \mathcal{L}}{\partial q_i}\delta q_i + \sum_{i,j}\frac{1}{2}\frac{\partial^2 \mathcal{L}}{\partial q_i \partial q_j}\delta q_i \delta q_j + \frac{\partial \mathcal{L}}{\partial t}\delta t + o(\delta t) \qquad (1)$$

Taking expectations of (1) we get (using the notation that <…> means expectation):

$$< \delta\mathcal{L} >= \sum_i \frac{\partial \mathcal{L}}{\partial q_i} < \delta q_i > + \sum_{i,j}\frac{1}{2}\frac{\partial^2 \mathcal{L}}{\partial q_i \partial q_j} < \delta q_i \delta q_j > + \frac{\partial \mathcal{L}}{\partial t}\delta t + o(\delta t) \qquad (2)$$

Note that $< \delta q_i \delta q_j >$ is quite different from $< \delta q_i >< \delta q_j >$.

Rearranging (2) we get:

$$\frac{\partial \mathcal{L}}{\partial t}\delta t =< \delta\mathcal{L} > - \sum_i \frac{\partial \mathcal{L}}{\partial q_i} < \delta q_i > - \sum_{i,j}\frac{1}{2}\frac{\partial^2 \mathcal{L}}{\partial q_i \partial q_j} < \delta q_i \delta q_j > + o(\delta t) \qquad (3)$$

We now substitute (3) into (1) and gather together some terms to arrive at the main result of the formalism:

$$\boldsymbol{\delta\mathcal{L} - < \delta\mathcal{L} >= \sum_i \frac{\partial \mathcal{L}}{\partial q_i}(\delta q_i - < \delta q_i >) + \frac{1}{2}\sum_{i,j}\frac{\partial^2 \mathcal{L}}{\partial q_i \partial q_j}(\delta q_i \delta q_j - < \delta q_i \delta q_j >) + o(\delta t)} \qquad (\wp)$$

---

[20] There are loose similarities with the notation and methods in this note to Lagrange's reformulation of Newton's laws but it would be wrong to expect the analogy to stretch too far

[21] If t changes by 1 week then the change in t is completely determined as 1 week!

Note that the term in $\frac{\partial \mathcal{L}}{\partial t}$ falls out. Assuming that our expectation of P&L is zero the LHS is the contribution to actual P&L and the RHS is its attribution.

In many situations there is a risk neutral measure whereby we have the following identities:

$<\delta\mathcal{L}>= r\mathcal{L}\delta t, \quad <\delta q_i>= rq_i\delta t, \quad <\delta q_i\delta q_i>= \sigma_i^2 q_i^2\delta t, \quad <\delta q_i\delta q_j> = \rho_{ij}\sigma_i\sigma_j q_i q_j\delta t$

where r is the risk free rate,

σ_i is the standard deviation of q_i and

ρ_{ij} is the correlation between q_i and q_j

in which case **(℘)** becomes:

$\delta\mathcal{L} - r\mathcal{L}\delta t =$

$$\sum_i \left(\frac{\partial \mathcal{L}}{\partial q_i}(\delta q_i - rq_i\delta t) + \frac{1}{2}\frac{\partial^2 \mathcal{L}}{\partial {q_i}^2}\left((\delta q_i)^2 - \sigma_i^2 q_i^2\delta t\right)\right) + \sum_{i>j}\frac{\partial^2 \mathcal{L}}{\partial q_i\partial q_j}\left(\delta q_i\delta q_j - \rho_{ij}\sigma_i\sigma_j q_i q_j\delta t\right) + o(\delta t)$$

where we have used the fact that $\frac{\partial^2 \mathcal{L}}{\partial q_i\partial q_j} = \frac{\partial^2 \mathcal{L}}{\partial q_j\partial q_i}$

Example 1:

As a simple example of the application of ℘ , let us consider a holding of stock *S*.

Its Lagrangian is as follows:

$S = Nq$ (4)

Where N is the stockholding and *q* is the stock price. Note that there is no explicit *t* dependence and that in general an object does not necessarily have a dependence on all the co-ordinates of the configuration space.

We have

$$\frac{\partial S}{\partial q} = N\,;\ \frac{\partial^2 S}{\partial q^2} = 0; < \delta S >= rS\delta t\ and < \delta q >= rq\delta t$$

Thus we get from ℘

$\delta S - rS\delta t = N(\delta q - rq\delta t)$

and substituting from (4) we get:

$\delta S = N\delta q$ which, of course, is exactly what we expect.

Example 2:

In this example we consider the Black Scholes Option value *(V)* with one degree of freedom (co-ordinate).

This time the co-ordinate *q* is the actual value of the underlying and we have the following:

$$\frac{\partial V}{\partial q} = \Delta\,;\ \frac{\partial^2 V}{\partial q^2} = \Gamma; < \delta q >= rq\delta t\ and < \delta q \cdot \delta q >= \sigma^2 q^2\delta t$$

Thus from $\wp$:

$$\delta V = rV\delta t + \Delta(\delta q - rq\delta t) + \frac{\Gamma}{2}((\delta q)^2 - \sigma^2 q^2 \delta t)$$

Where we have used the risk neutral assumption that $< \delta q >= rq\delta t;\ < \delta V >= rV\delta t$

Bibliography

Options, Futures, and other Derivatives: John C. Hull

Financial Calculus: Martin Baxter and Andrew Rennie