Pricing and hedging of variable annuities

Variable annuity products are unit-linked investments with some form of guarantee, traditionally sold by insurers or banks into the retirement and investment market. Pricing VAs is similar to pricing long-dated financial derivatives on a basket of assets. In this paper we explore the pricing, sensitivity and risk management of a specific GMWB contract

BY DANIEL BLAMONT AND PRETTY SAGOO

BY FAR THE LARGEST VA market in the world is the United States, where variable annuities are widely used in retirement planning. In the absence of compulsory annuity purchase requirements on retirement such as in the UK, VAs are a common and long-established retirement investment with some tax advantages.

Variable annuities are also referred to as accumulator products or as GMxBs where, in the latter, the 'x' describes the nature of the guarantee in the product, commonly known as the rider. A typical 'Guaranteed Minimum x Benefit' might refer to Accumulation (GMAB), Income (GMIB), Death (GMDB) or Withdrawal (GMWB).

In the US, which has the best established VA market, the most common VA product is by far the GMWB. Approximately 78% of VA sales in the US in the first half of 2005 contained a GMWB feature[1]. The focus of this paper is therefore a dollar-denominated GMWB policy. More recently, the GMWB for life, introduced in 2004, has proven extremely popular.

When a GMWB is purchased the initial capital is invested in a sub-account at the holder's discretion. The holder can withdraw guaranteed periodic amounts up to the value of the initial capital. The GMWB terminates once the initial capital has been withdrawn; any remaining funds in the sub-account are returned to the policyholder at maturity. So the sub-account value fluctuates with movements in the underlying assets and decreases with withdrawals.

Therefore, a GMWB effectively combines:

- an annuity, in the form of guaranteed periodic withdrawals
- a call option on the underlying residual sub-account at maturity

The policyholders are the owners of the sub-account. As such the insurance company is selling protection on an account that is largely managed by the policyholder. Typically, these products offer clients some restricted choice of investment funds for the sub-account with some limited ability to switch fund allocations during the lifetime of the policy. An asset allocation split of 60/40 between equities and fixed income is quite common.

The potential uncertainty arising from fund switching or varying amounts of equity in the underlying asset allocation should imply high fees. However, finite-life GMWBs typically carry a fee of 40-60 basis points per annum, traditionally charged as a percentage of the sub-account value. We will see in the next section how fees tend to understate the financial cost of the protection provided.

Table 1: Annual charges

	Typical annual charge
GMDB	15–35bp
GMWB	40–60bp
GMAB	30–75bp
GMIB	50–75bp

Source: Deutsche Bank

Insurance companies compete against each other not only on fees but also on the features of the product. Nevertheless, the typical GMWB carries the following features:

- fee charged to the sub-account
- some fixed maximum annual withdrawal per $100 invested initially, typically $7
- withdrawals above the limit are allowed but the investor pays a penalty and may give up some of the guarantee

Figures 1 and 2 illustrate the workings of a typical GMWB for two given sub-account scenarios: a bad one and a good one.

1 VARDS (Variable Annuity Research and Data Service)

Figure 1. GMWB – good path

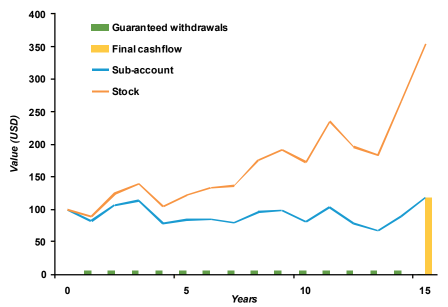

Source: Deutsche Bank

Figure 2. GMWB – bad path

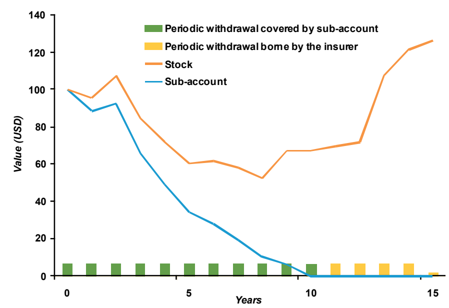

Source: Deutsche Bank

- 5% flat yield curve
- lognormal (Black-Scholes) underlying asset with 20% volatility
- $7 yearly withdrawal
- $100 initial investment/guarantee/stock price
- no fees and no penalties

In other words, the underlying asset satisfies the following standard stochastic differential equation under the risk neutral measure:

$$\frac{dS}{S} = rdt + \sigma W_t \quad ,$$

In a good scenario (Fig. 1), when the underlying assets perform well, the yearly withdrawal of $7 will be at least partially offset by capital appreciation. The sub-account therefore covers the withdrawals over the life of the GMWB and even has a positive value at maturity. The investor benefits from the periodic cashflows, as well as the terminal value of the sub-account. The insurer receives the annual fees and has no shortfalls to cover.

In a bad scenario (Fig. 2), there comes a point where the sub-account value cannot cover the $7 withdrawal. At that point the sub-account is liquidated and the balance of $7 is covered by the insurer. In addition, until the GMWB maturity, the insurer will have to cover the $7 periodic withdrawals. In this instance the investor only receives a yearly $7 cashflow with no upside. The insurer provides these cashflows when the sub-account reaches $0 until maturity.

Pricing

The pricing of GMWBs and VA policies in general is similar to pricing long-maturity exotic financial derivatives. VAs have the following features in common with exotic derivatives:

- A basket of underliers (equity, fixed income and others), ie. hybrid features
- Ratchets (lookbacks), or roll-up features

The biggest difference between VA policies and most financial derivatives is that the former combine financial risks with insurance risks such as surrender, longevity and mortality. In this sense, they are complicated products to price, and even more complicated to hedge precisely.

In the case of our GMWB, the risk that more, or less, people than expected surrender their policies, and the timings of these surrenders, can have a significant effect on pricing. We will explore the effects of insurance risks on pricing and risk management in a further paper.

In general, at least for market risks, similar rules apply as would to standard derivatives; a longer maturity product would tend to cost more and will be sensitive to the level and volatility of the underliers.

Deterministic rates

We initially assume deterministic rates. We will relax this assumption later on and investigate its impact. Throughout this section the following will be assumed as a base case:

where:

W_t is a Weiner process,
r is the interest rate, set at 5%,
and σ is the volatility, set at 20%.

Two methods of pricing are widely used: Monte Carlo (MC) and Numerical Partial Differential Equations (PDE). The MC method has the advantage of allowing a great number of variables to be stochastic and any payoff to be priced. In addition it will lead to a distribution of outcomes, which proves useful when doing scenario analysis. The PDE methods can be more accurate and further can be used to consider optimal withdrawals, unlike MC (Milevsky and Salisbury, 2004). Initially we will use the MC method for pricing purposes.

Following put-call parity arguments, the value of the guaranteed periodic withdrawals and the call option on the sub-account (residual sub-account at maturity) should equate to the sub-account plus the cost of insurance (getting the guarantee when the sub-account is zero, ie., a put option):

Call option + Guaranteed withdrawals
=
Sub account + Insurance cost

Figures 3 and 4 show the main variables along the average path (average over the simulations). Over the average path the present value of the yearly $7 withdrawals is $71.3. The present value of the expected residual sub-account value is $32.7. This gives a total $104.0 value to the investor.

Figure 3. Call + guarantee value on the average path

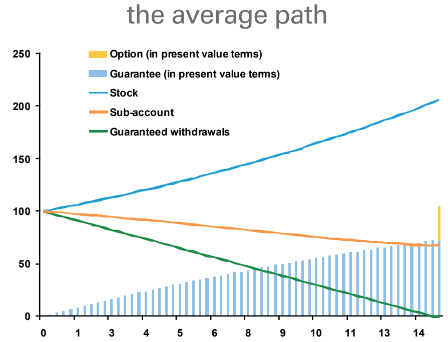

Source: Deutsche Bank

Figure 4. Cost of insurance on the average path

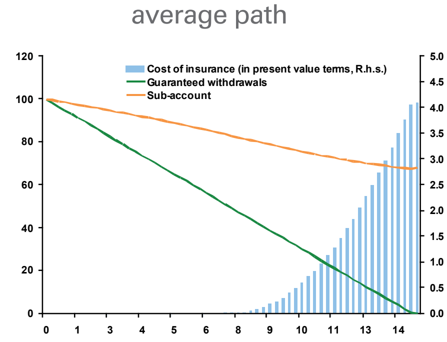

Source: Deutsche Bank

Table 2: Pricing for various maximum withdrawal levels				
Withdrawal	Option	Guarantee	Package	Insurance
5	40.27	63.08	103.35	3.35
7	32.73	71.31	104.05	4.05
10	26.03	78.53	104.55	4.55
15	19.93	84.86	104.79	4.79
20	16.60	88.29	104.89	4.89

Source: Deutsche Bank

This therefore implies that the cost to the insurer is \$4.0. In theory, fees on the sub-account should cover that cost so that the expected value to the investor is less than \$100, implying a negative cost to the insurer.

Table 2 shows the value of the different GMWB components for various allowed withdrawal levels. Since we assume the investor withdraws the maximum allowed, withdrawals are another way of defining the product maturity. A withdrawal of w per year implies a maturity of $100/w$ years.

As such, high withdrawals reduce the option value and increase the guarantee value. The guarantee value is obviously increased as the time value loss is smaller since the \$100 nominal value is recovered quicker. The option value is reduced as there is less time for the sub-account to grow.

More importantly, the higher the withdrawal allowed by the GMWB, the higher the insurance cost. The act of withdrawing a high proportion of the sub-account means it will converge quicker to zero and have less time to recover from a bad underlying asset performance. As such the insurer is more likely to have to provide the guarantee.

So what should be the fair value cost of investing in GMWBs? We have seen that it will depend on the withdrawal level. It will also depend on the volatility of the underlying sub-account assets. The more volatile the assets, the more valuable the option. Table 3 shows some fair value fees, assuming the fee is deducted as a portion of the sub-account (effectively equivalent to a dividend yield on the underlying asset).

Table 3: Fair value fee (bp per annum)			
	Volatility		
Withdrawal	10%	20%	30%
4	1.0	16.5	46.0
7	5.0	52.0	132.5
10	11.5	97.0	226.0
15	27.0	165.0	367.0

Source: Deutsche Bank

As expected, the fair value fee grows with withdrawals as it increases the value of the insurance. It also reduces the maturity of the GMWB, shortening the period over which the fee can be earned. This aspect is only important at the product design phase.

The fair value fee grows exponentially with volatility and therefore highlights how critical the volatility of the underlying asset is. It is therefore important for insurers to consider the likely volatility of the assets chosen by the investor.

The results imply that some insurers may have underpriced GMWBs, should investors elect volatile investments. As shown earlier, typical fees range from 40 basis points to 60 basis points. This apparent underpricing may have arisen from changes in product design or competitive pressures.

The underpricing becomes more apparent when we consider that the fair value fees that we have deduced in Table 3 are just the theoretical cost of hedging the guarantee. In reality, the fees charged to the policyholder need to include cost of capital for the policywriter, frictional costs arising from the basis risk of the theoretical versus actual hedge, a profit margin for the policywriter and any fees needed to cover insurance risks that have not been accurately captured in the pricing models used.

On this latter point, the insurance risks in the VA products being offered, such as lapse or mortality, can have a large impact on pricing. In the case of lapse, or policyholder behaviour in general, in the absence of a rigorous pricing methodology for the actual risk, VA writers need to price conservatively and thus need to assume that policyholders behave completely rationally (ie. they always exercise in-the-money options). Traditionally, this assumption has not been made and therefore has added to the degree of mispricing that has occurred in this market.

Furthermore, many VAs were historically priced under real-world assumptions (P-measure), including risk premia on assets. For hedging and fair valuation purposes the risk neutral Q-measure is the one that matters, under which all assets earn the risk-free rate. The implication is that indeed, on average, in the real world, fees might cover for the expected insurance cost. However, in practice there is only one realisation of asset returns. In that case, only if fees are used to implement the appropriate hedging will they cover for the expected insurance costs.

We now investigate the greeks for the typical GMWB. For that purpose numerical PDEs are a more appropriate tool. Under the PDE methodology, the GMWB is priced recursively starting at maturity. We can therefore derive a surface of GMWB value versus time and sub-account value.

From this we can derive the insurance value, as well as its delta and gamma. We calculate the basic greeks with respect to the total account value, for demonstration's sake. However, from a risk management practitioner's point of view, risk needs to be analysed with respect to the level of component risk factors in the account basket, ie. interest rates and equities (rho and delta respectively), and their volatilities (vega).

The insurance cost (Fig. 5) exhibits a put-like profile. The main difference with a standard put is that the forward path is one where the underlying drops in value, because of withdrawals, rather than growing at the risk-free rate.

Figure 5. GMWB insurance cost

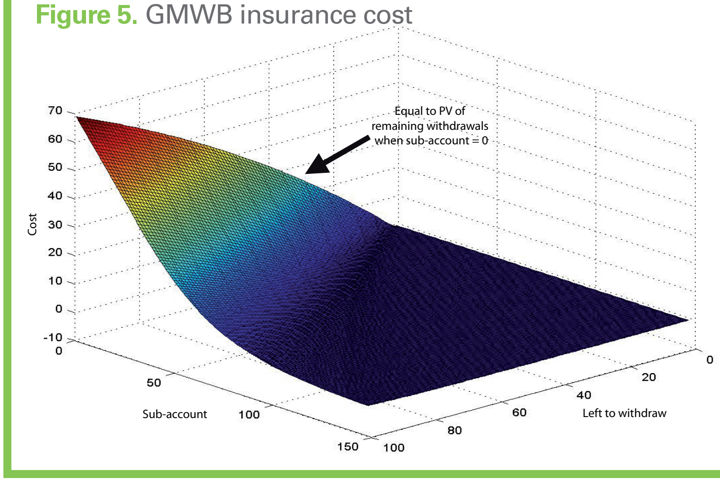

Source: Deutsche Bank

Figure 6. GMWB insurance delta

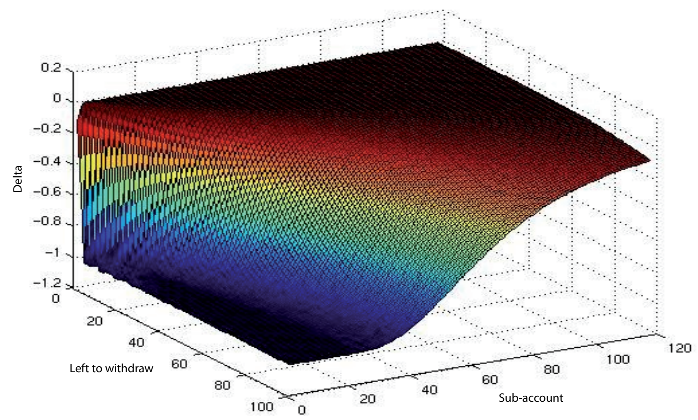

Source: Deutsche Bank

Figure 7. GMWB insurance gamma

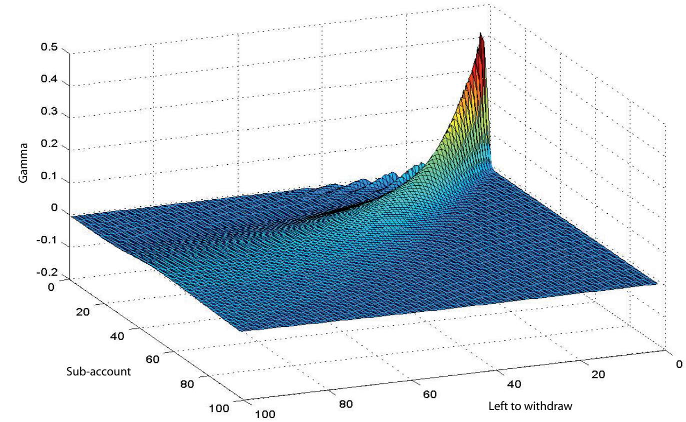

Source: Deutsche Bank

Similar to a vanilla option, the delta (Fig. 6) exhibits an S-shaped profile and the gamma (Fig. 7) peaks around the 'at-the-money' level and increases as we get closer to maturity.

Stochastic rates

Until now we have only considered deterministic rates and a flat yield curve. Instead, we can consider stochastic rates. We do so here, using a single factor Vasicek model, calibrated on the US yield curve (as of August 31, 2008). In this model, the short rate dynamics under the risk-neutral measure Q are given by $dr(t) = k[\theta - r(t)]dt + \sigma dW(t)$, where k, θ, and the volatility σ are constants and W(t) is a Brownian motion. As we introduce another stochastic variable, we only consider Monte Carlo pricing.

Under this set-up our results are qualitatively similar to the ones obtained earlier. In fact, the most important factor now is the correlation between rates and the underlying assets.

When quoting on hybrid products, traders usually calibrate their bid-ask spread based on rates-equity correlations of -50% and +50%. We calculate the GMWB insurance cost based on these correlations as well as 0%.

As shown in Figure 8, the insurance cost generally increases with the allowed withdrawal amount (or as maturity decreases), just like under deterministic rates. The insurance cost increases with correlation, similarly to a standard put. With positive correlation, a put option will tend to pay off when rates are low. The payoff will be discounted at a lower rate, giving a higher value. The opposite happens with negative correlation.

Figure 8. Insurance sensitivity to rates-equity correlation

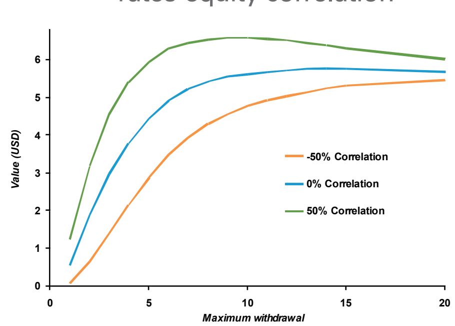

Source: Deutsche Bank

Hedging

Local and international regulators require insurance companies to appropriately measure financial risks underlying their business. As we move towards Solvency II there is increased focus on market consistent risk measures including tail risk, such as value-at-risk measures. Hedging can have a dramatic impact on such tail-risk measures.

Fees levied on GMWB are set to cover the cost of protection on average. In the real world there is only one realisation of asset performances. It is therefore important to invest fees in assets that will replicate the behaviour of the guarantee being sold.

Hedging versus reserving

For regulatory and economic purposes, the rationale for hedging or reserving is to be able to withstand a significant market shock. The relevant risk level is typically the 0.5% percentile over a one-year horizon.

Assuming normal distributions, this requires covering against a 2.58 standard deviation shock. In other words, if assets are kept as a reserve against deterministic shocks, solvency is only guaranteed at the confidence level considered. Only reserves equal to assets can ensure solvency with 100% confidence.

On the other hand, one can hedge the tail risk via put options. In this case the tail risk will be entirely covered with 100% confidence. The burden of providing protection, including model and jump risk will fall on the option writer.

The advantage of hedging is that capital allocated to hedges can be many times smaller than reserving. In Table 4 we demonstrate this for the situation where the VA writer is short a put only, such as for a simple GMAB on a basket of interest rates and equities. In this case, the hedge cost is the cost of the put whereas the cost of reserving for the risk is equivalent to holding

Table 4: Hedging versus reserving

	Hedge	Reserve
Cost[2]	$0.4 \times \sigma_A \times \sqrt{T}$	$\sim 2.6 \times \sigma_A \times \sqrt{T}$
Advantage	Total transfer of risk	Investment freedom of free assets
Disadvantage	Negative carry asset (time value decay) Volatility risk premium cost	Still non-zero probability of default Cost of capital and de-leveraging of balance sheet in distress periods

Source: Deutsche Bank

2 The first formula is an approximation of the Black-Scholes formula where the option is struck at the money forward ($K = Se^{rT}$). The Black-Scholes formula collapses to Put = Call = S [N(d1) – N(d2)] with d1 = –d2 = ½ σ sqrt(T). A Taylor expansion of the function N(.) at d1 and d2 yields to the following approximation: Put = Call ~ 0.4 S σ sqrt(T). This cost estimate ignores transaction charges.

assets sufficient to absorb a shock down to the reserve level. Assuming a stress equivalent to a Solvency II stress level of 99.5%, this requires assets of roughly 6 times the cost of hedging.

The decision to hedge or not should therefore be based on the cost of hedging versus the opportunity cost on reserve capital (Figure 9).

Static versus active hedging

When it comes to hedging, a wide spectrum of solutions is available, from static hedging on one side to active hedging on the other.

In the case of GMWBs, the perfect hedge is a put option on the sub-account value. Given that the sub-account value is affected by relative changes (performance of the underlying) as well as jumps (withdrawals) it cannot be hedged with standard instruments. Milevsky and Salisbury (2004) claim that the value of the GMWB is equivalent to a Quanto Asian put on an underlying which is the inverse of the sub-account price. This still keeps the sub-account as underlying and introduces an exotic option which would still have to be managed using plain vanilla options. In other words, there will always be a basis risk between the hedge and the required hedge.

A perfect hedge is likely to be prohibitively expensive given the complexity of the underlying sub-account process. A static hedge is more likely to be a portfolio of options chosen to best replicate the insurance cost. A basis risk will remain. A dynamic hedge should reduce the basis risk but will introduce transaction costs as well as requiring ongoing risk management.

To illustrate the pros and cons of both strategies we will reconsider the GMWB with deterministic rates. For the dynamic hedge we will assume a monthly rebalancing. Standard delta-hedging is undertaken giving rise to a basis risk given that only the first derivative with respect to the underlying is considered. Other greeks and higher orders are ignored. The reader should note that delta hedging alone, albeit on a more frequent basis, has been the hedging approach adopted by many US insurers who have sold VA products. The technique, in the absence of hedging of other and higher order greeks, has proven to be insufficient in stress scenarios, particularly where large sudden shocks occur in market levels and volatility.

In the static case we can consider a range of potential solutions:

- Single put
- Portfolio of puts with different strikes and maturities
- Put spreads, designed to only provide a $7 payoff at option maturity, in accordance with withdrawals

The allocation to various instruments is performed by minimising the basis risk at maturity between the insurance cost and the hedge payoff. Effectively we perform a constrained regression (positive weights, no intercept) to replicate the insurance risk with our universe of hedges. The R^2 of the regression is the metric we use to assess the effectiveness of the hedge. As a result, the optimisation of a static hedge still introduces model risk.

Initially we consider the single put (Fig. 10). For that we look at a selection of puts with maturities ranging from 5 to 14 years with strikes between 80% and 110% of spot. We find that at-the-money spot puts work best and the 7-year maturity is optimal, giving an R^2 of 81.1%. Longer-dated puts are less effective, given the path dependency of the insurance cost and the greater impact of bad performance of the underlying in earlier years.

Figure 9. NAV profile at maturity

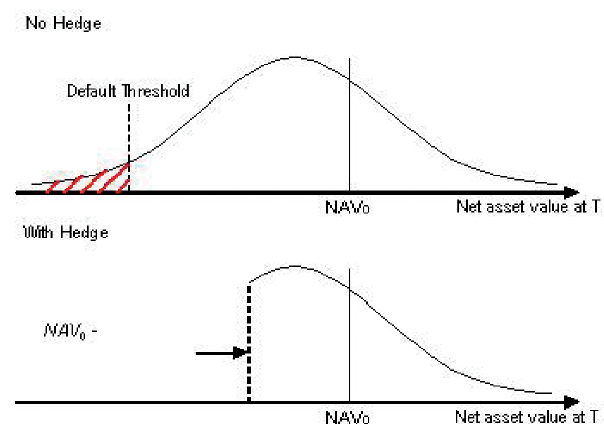

Source: Deutsche Bank

If we consider a portfolio of at-the-money puts (5- to 14-year maturities), we can increase the goodness of fit to 93.7% (Fig. 11). Under this setup most of the weight is attributed to the 5-year option. In general, longer-dated puts attract smaller weights, in line with the decreasing R^2 at individual put level. A portfolio of puts also has the advantage of not having the hedge concentrated in a single instrument.

The portfolio of put spreads performs marginally better than the single put hedge (R^2 of 86.4% versus 81.3%), but worse than the portfolio of standard puts (R^2 of 93.7%). The rational for put spread was only to get a payoff equal to the required amount, ie. $7 to cover for the withdrawals in any particular year. However, the underlying of the put is different from the sub-account. In fact the binary nature of the payoff of the put spreads introduces further basis risk.

But even the best static hedge underperforms the dynamic hedge (R^2 of 96.6%). Figures 12 and 13 show how the dynamic hedge always leads to a small basis risk. The static hedge sometimes leads to no basis risk when the underlying asset performance is good (no insurance has to be provided) but leads to a basis risk when the underlying performance is poor.

This is further highlighted in Figure 13 showing a much tighter distribution of the basis risk for the delta hedging than for the portfolio of puts. We also see that the portfolio of puts significantly reduces downside basis risk relative to a single put.

It is important to remember that in this example we have ignored insurance risks, in particular, policyholder behaviour. The dynamic approach provides

Figure 10. Single puts: goodness of fit

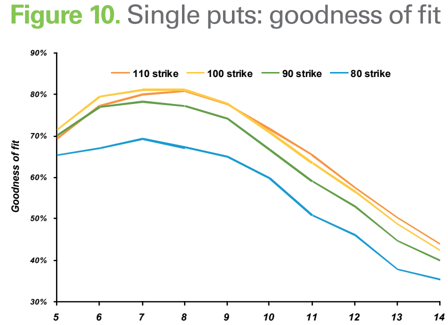

Source: Deutsche Bank

Figure 11. Portfolio of puts: individual weights

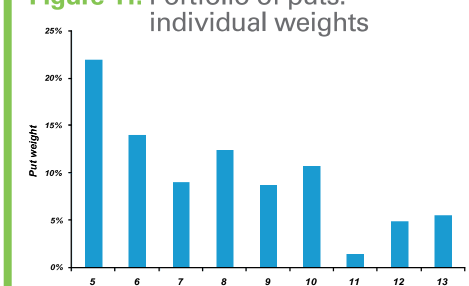

Source: Deutsche Bank

Figure 12. Basis risk and final equity value

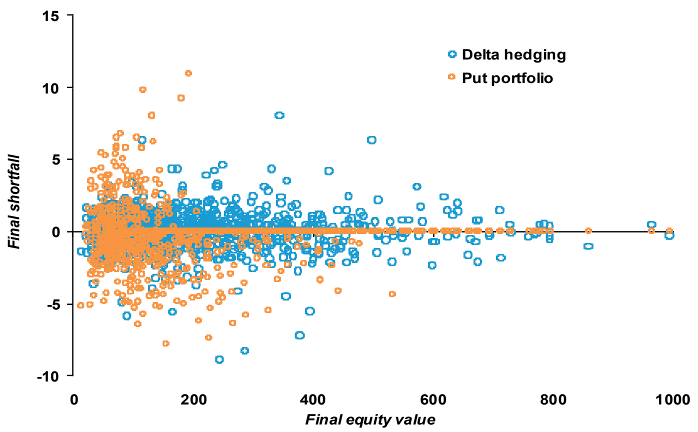

Source: Deutsche Bank

Figure 13. Basis risk and final equity value

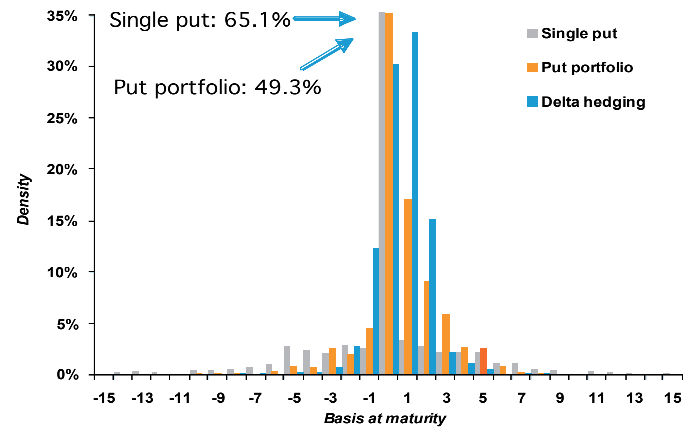

Source: Deutsche Bank

flexibility in managing insurance risks through the constant re-appraisal of the hedge as insurance risks materialise. The dynamic approach would therefore be superior in a real-world setting.

Alternatively, one could hedge semi-statically. In other words: dynamically managing a static hedge because, although the dynamic hedge appears optimal, in practice there may be hurdles to implementing it. First, there is model risk and in particular the possibility of jumps. This is nearly impossible to delta hedge. With a static hedge, the burden of replication in case of jumps (and the options management in general) would fall on the option writer, while risk monitoring would remain in-house with the VA seller. Furthermore, carrying out a dynamic hedge requires appropriate resources which not all insurers have or could afford.

It is worth mentioning one way in which VA sellers can remove most of the risks arising from the VA book. This involves transacting a reinsurance type contract, with either a bank or reinsurance company. The last remaining risk then becomes counterparty risk – which can further be limited through collateralisation. In current markets, we see the ability to transact such contracts as being increasingly limited, with many reinsurers reassessing their appetite for such deals.

For those who choose to retain insurance risks, in the absence of a (liquid) market for hedging these risks, they normally remain with the VA seller. Nevertheless, they can be mitigated through product design and through reserving explicitly for risks that cannot currently be hedged through the capital markets.

Conclusion

While this article has looked in some detail at pricing and hedging of one particular type of VA policy, it should have become clear that pricing and hedging these products 'correctly' is a very product-specific and potentially complicated exercise. We considered a very simple product – the addition of 'attractive features' such as ratchets or lookbacks further complicates the pricing and hedging process of VA policies.

In the context of an insurer taking such products to market for the first time, as is the case for many European insurers, the list of practical considerations expands dramatically:

- Can current distribution channels fairly market the product?
- Is product design fully optimised for risk management?
- Does intellectual and technical pricing capability exist in the firm?
- What reserves are required?
- What risk management techniques should be adopted? Dynamic or static hedging?
- Are new systems required for risk monitoring and/or hedging?
- Is reinsurance a possibility?

In the US and Japan, where VA products have been around for a while, VA sellers have already dealt with the above considerations, even if, in our opinion, charges for VA products tended not to cover fully the implied costs. In these markets, VA sellers have entered a new phase of risk management of these products such that a re-appraisal of the risk versus charges is under way. This has led VA sellers to consider outsourcing hedging to third parties more naturally equipped or willing to deal in dynamic market and insurance risks. L&P

Daniel Blamont and Pretty Sagoo are vice-presidents of the Strategic Solutions Group and Global Pensions Strategy Group respectively, at Deutsche Bank
Email: daniel.blamont@db.com, pretty.sagoo@db.com
The opinions or recommendations expressed in this article are those of the authors and are not representative of Deutsche Bank as a whole

References

Milevski, M. and Salisbury, T.
Static and Dynamic Valuation of Guaranteed Minimum Withdrawal Benefits
August 2004

Merton, R. and Perold, A.
Theory of Risk Capital in Financial Firms
Journal of Applied Corporate Finance, 6(3) (Fall 1993, 16-32)

Bauer, D., Kling, A. and Russ, J.
A Universal Pricing Framework for Guaranteed Minimum Benefits in Variable Annuities
Working Paper, Ulm University