Understanding the Behaviour and Hedging of Segregated Funds Offering the Reset Feature

by

H. Windcliff
Department of Computer Science
University of Waterloo
Waterloo ON
Canada N2L 3G1
hawindcliff@elora.math.uwaterloo.ca

M.K. Le Roux
Sun Life Financial
225 King Street West
Toronto ON
Canada M5V 3C5
martin.le.roux@sunlife.com

P.A. Forsyth
Department of Computer Science
University of Waterloo
Waterloo ON
Canada N2L 3G1
paforsyt@elora.math.uwaterloo.ca

and

K.R. Vetzal
Centre for Advanced Studies in Finance
University of Waterloo
Waterloo ON
Canada N2L 3G1
kvetzal@watarts.uwaterloo.ca

Acknowledgement: This work was supported by the Natural Sciences and Engineering Research Council of Canada, the Social Sciences and Humanities Research Council of Canada, Royal Bank of Canada, and Sun Life Financial.

Abstract

Segregated funds have been an extremely popular Canadian investment vehicle during the past few years. These instruments provide long term maturity guarantees and often include extremely complex option features. One heavily debated feature is the reset feature; the ability to lock in market gains. In typical cases, investors have the ability to do this two or four times per year. The valuation of this embedded optionality has been controversial. Recently, regulators have announced that firms offering these products will be subject to new capital requirements. In this paper we discuss the effects of market parameters, such as volatility and interest rates, on the cost of providing a segregated fund guarantee. We also demonstrate how the level of optimality in investors' use of the reset feature affects the cost of providing such a guarantee. For each scenario, we provide the appropriate management expense ratio (M.E.R.) which should be charged as well as demonstrating the current liability using a given fixed M.E.R. Further, we explore ways of modifying standard contracts so as to reduce the required hedging costs. We also take a closer look at some intuitive reasons why the reset feature requires such a dramatic increase in the hedging costs. Finally, we present an approximate method for handling the reset feature which can be computed very efficiently. This method provides accurate results when the correct proportional fee is being charged.

1 Introduction

In terms of total sales volume, segregated funds have been one of the most successful Canadian financial products over the past several years. Essentially, these contracts are mutual fund investments, augmented with additional features provided by insurers. These extra features include a maturity guarantee (typically for a period of ten years) and mortality benefits. In addition, in many cases investors have the ability to switch from one underlying fund to another (while maintaining their guarantee levels) and also to reset their guarantee. This reset feature has sparked a great deal of controversy and debate. It allows the investor to lock in market gains at any time, up to a maximum, for example, of two or four times per year. When the reset feature is used, the maturity date of the guarantee is changed to ten years from the date of the reset.

In light of the stricter capital requirements recently imposed by OSFI[1], some companies have begun selling more restricted versions of these contracts. Other companies have discontinued the sale of these contracts altogether, deeming the capital requirements to be too onerous. However, large volumes of contracts already sold remain outstanding, and companies need to manage their risk exposure, or at least allocate sufficient capital and reserves.[2]

The purpose of this paper is to shed some light on the intuition behind the valuation of the reset option embedded in segregated funds. The results we present are derived from a financial option pricing model based on the numerical solution of a set of linear complementarity problems, described in detail in Windcliff et al. (2001). As our focus in this paper is exclusively on results, we shall not describe the model in detail here. Readers who are interested in more of the technical aspects should consult Windcliff et al. (2001).

We find that the option features provided by these contracts are so complex that intuition often fails. In particular we address the following questions:

- Why does the reset feature seem to be so valuable?
- Can we modify these contracts so as to reduce the hedging costs?

Note that our "heuristic" and "intuitive" descriptions of issues related to these contracts are not meant to be taken as rigorously correct. Instead, they are intended to illustrate some of the subtle points which are very elusive to track down, even though in some cases they will seem obvious in retrospect. We emphasize, however, that our results are based on a rigorous model.

We also provide concrete demonstrations of the sensitivity of the value to some common modeling input parameters. These include variations in the volatility of the underlying mutual fund and

[1]The Office of the Superintendent of Financial Institutions (OSFI) is a Canadian regulatory agency.

[2]According to the consulting firm Investor Economics, the total amount of individual segregated fund contracts issued by Canadian insurers was \$45 billion (CDN) as of August 2000, up from \$9 billion at the end of 1995. 47% of these assets had a maturity guarantee of 100%. Note that although not every insurer provides optionality such as reset provisions, many of them either do or did at one time.

changes in the risk free interest rate. Another input parameter which is important in the valuation and hedging of these contracts is the level of optimality displayed by the investors, both in terms of their use of the reset option and in terms of their lapsation behaviour.

In this paper, we use two different measures of the cost of providing the guarantee portion of these contracts.

Appropriate M.E.R.: Given a particular contract and market conditions, we would like to determine the appropriate proportional fee that the customer should pay in order to cover the cost of providing the guarantee portion of the segregated fund contract. We assume that the insurer charges a proportionate fee, r_e, to cover its risk management expenses, over and above the proportionate fees it charges to cover investment management expenses, administration costs, distribution costs, and the cost of capital.

Residual cost, net of future incoming fees: If the correct proportionate fee is being charged, then, by definition, the contract has zero net value to the insurer at inception (treating the cost of capital as an expense). However, if a different fee is being charged, or if fund market values have changed, then we would like to determine the value to the insurer of a contract given a particular proportional fee.

Throughout this paper we will see a central theme. If one adopts a no-arbitrage perspective (as in modern financial option valuation), in many cases these contracts appear to be significantly underpriced, in the sense that the current deferred fees being charged are insufficient to establish a dynamic hedge for providing the guarantee. This is particularly true for cases where the underlying fund has a relatively high volatility. This finding might raise concerns at institutions writing such contracts.

2 General Description of Model and Key Assumptions

In this section we describe some of the main aspects and assumptions of the approach we use for valuing the guarantee portion of a segregated fund. We use modern financial option pricing theory to solve for the *no-arbitrage* value of the embedded options (both the initial guarantee, which can be viewed as a simple put option, and the reset option). The basic idea involves the construction of a dynamic hedging portfolio to replicate the payoffs of the option components. As this portfolio is *self-financing*, the entire cost of setting it up is incurred at the outset. Since the hedging portfolio produces the same payoffs as the embedded options, the no-arbitrage principle (i.e. any two assets with identical payoffs in future circumstances must have identical values today) implies that the value of the options is equal to the initial cost of the hedging portfolio.[3]

[3] Extensive treatments of financial option valuation can be found in texts such as Hull (2000) or Wilmott (1998).

As a general observation, note that we use a numerical partial differential equation (PDE) method which we solve backwards in time. This allows us to develop a rigorous treatment of optimal resets. By solving backwards in time, we are always able to determine the required cost to hedge the remaining portion of the contract. In contrast, techniques such as Monte Carlo simulations are solved forward in time. Although there has been a lot of recent progress in this area (see Boyle et al., 1999, for example), Monte Carlo techniques cannot handle resets with the same level of rigor. As noted above, the low level details of our approach are described in Windcliff et al. (2001) and will not be repeated here.

However, it is important that we explicitly state some of the key assumptions made in using our approach. First, we assume that the dynamic hedging portfolio can be formed using a basket of stocks which are contained in the underlying mutual fund, but not the mutual fund itself. In other words, hedging is based on an underlying asset where no M.E.R. is charged. This also implies that there is no *basis risk* introduced by the hedging. Since the underlying mutual fund can be actively managed, this may be very difficult to achieve in practice. Thus the results computed here are a lower bound to the no-arbitrage value, and additional reserves will be required for this basis risk. Second, we assume that markets are frictionless. In theory, the dynamic hedge is adjusted continuously through time at zero transactions costs. In practice one would use a discretely adjusted hedging portfolio, implying that the payoffs of embedded options will not be perfectly replicated by the hedging portfolio. Moreover, non-zero transactions costs will also increase the cost of hedging and thus the value of the options.[4] Finally, we also assume that the volatility of the underlying mutual fund and the risk free rate of interest are deterministic (known) processes. In this paper, for reference, they are taken to be constant (i.e. the risk free rate is constant, and the underlying mutual fund evolves according to geometric Brownian motion). In other words, we adopt the standard Black-Scholes framework. This has well-known limitations, and many practitioners compensate by using the volatility parameter as a kind of "fudge factor". Nonetheless, this framework does provide a useful benchmark. We remark that there has been considerable research into more sophisticated option valuation models in the context of simple equity options. Examples include constant elasticity of variance models (Cox, 1975), stochastic volatility models (Heston, 1993), jump-diffusion models (Merton, 1976), models with stochastic interest rates (Rabinovitch, 1989), etc. Incorporating the additional complexity of these models into the segregated fund context presents an interesting possibility for future research.

3 Base Scenario

A summary of the important features of a prototypical segregated fund guarantee is given in Table 1. This is the base scenario for the valuations in this paper. It is our intention to show that even this

[4] See Boyle and Hardy (1997) for a discussion of these issues in the case of simple maturity guarantees, without reset provisions.

Investor profile Investor optimality Deterministic lapse rate Optimal (anti-selective) lapsing	50 year old female. 25% p.a. 5% p.a. Yes.
Initial investment Maturity term Resets	\$100 10 years, maximum expiry on investor's 80^{th} birthday. Two resets per year permitted until the investors 70^{th} birthday. Upon reset: Guarantee level = Asset level Maturity extended by 10 years.
M.E.R. D.S.C.	Total M.E.R. of 3.0% with 2.5% allocated to fund manager of underlying fund and $r_e = .5\%$ being used to fund maturity guarantee. Sliding scale from 5% in first year to 0% after 5 years in fund.
Volatility Interest rate	$\sigma = 17.5\%$ $r = 6\%$

TABLE 1: *Valuation assumptions for base scenario.*

relatively simple contract is very expensive (given our assumptions about the risk free rate and the level of volatility). Contracts which provide further optionality can be dramatically more expensive to hedge.

We will assume an initial investment of \$100. The investor initially receives a guarantee at this level which matures in 10 years. As is well known (see, e.g. Brennan and Schwartz, 1976; Boyle and Schwartz, 1977), this guarantee can be thought of as a ten year European put option with a strike price of \$100 (ignoring mortality provisions). However, additional complexity is introduced by the fact that the guarantee level can be reset by the investor up to two times per year. Upon reset, the guarantee level is set to the prevailing value of the fund, and the maturity date is extended to be 10 years from the reset date.

No initial fee is charged to enter into the contract. Instead, a total M.E.R. of 3.0% is charged. Of this, 2.5% covers investment management expenses and other expenses as described above and the remaining 50 basis points (b.p.) is used to compensate the insurer for providing the guarantee. Back end fees are charged upon early redemption. In this work, we use a sliding scale from 5% in the first year to 0% in the sixth and further years. It is assumed that the deferred sales charge (D.S.C.) goes entirely to the management of the underlying mutual fund. Consistent with standard practice, none of this fee is allocated for funding the guarantee portion of these contracts.

A standard mortality feature is provided. If the investor dies, the guarantee is provided immediately (at the time of death). In this work, we use mortality data for a Canadian female aged

50 years. In Windcliff et al. (2001), we show that the contribution of the mortality feature to the value of these contracts is minimal for this demographic type. Of course, for older investors with higher mortality rates, this may no longer be the case. The contract expires after the investor's 80^{th} birthday; i.e. the maximum duration of the contract is 30 years. The investor is not allowed to reset the guarantee level after her 70^{th} birthday. Again, we are trying to demonstrate that even relatively simple segregated fund guarantees with the reset feature can be very valuable. Some companies offer guarantees which allow the investor to reset the guarantee during the final years of the contract and permit maximum maturity dates which are more aggressive than the 80 year age limit we impose here.

The value of the guarantee will also depend upon market conditions. Here we assume that the guarantee is provided on a moderate volatility fund with $\sigma = 17.5\%$. We also assume a risk free interest rate of $r = 6\%$.

Finally, we need to make some assumptions regarding the behaviour of contract holders. First, the level of investor optimality, i.e. how *effectively* they use the reset feature, will have important effects on the cost of providing the guarantee portion of these contracts. For the base case, we assume a level of optimality of 25%. By this we mean that if it is optimal to reset, then over a one year interval 25% of the investors will do so.[5] Second, investors can lapse out of the contract. This can happen in two ways:

- For liquidity or other reasons, some investors will simply choose to close their accounts. We assume that 5% of the accounts are lapsed deterministically each year; i.e. independently of the value of the fund.

- In some situations, lapsing can be an optimal strategy for investors. For example, if a very high M.E.R. is being charged, it may be optimal for investors to lapse to avoid paying the fees. This is sometimes referred to as "anti-selective" lapsation. It is particularly likely if the value of the embedded options is fairly low (say, because all reset opportunities have been used and the underlying fund value is well above the guarantee), at least relative to the cost of paying the fees to remain in the contract.[6] For anti-selective lapsing, we use the same level of investor optimality as we use for triggering the reset feature, i.e. 25%.

Note that lapsation in and of itself cannot be assumed to be beneficial to the insurer when a deferred fee is charged. Lapsation reduces future fee income. As pointed out in Windcliff et al. (2001), if this happens in a situation where it is unlikely that any payoffs will ultimately have to be made, the insurer can be worse off.

[5]This interpretation assumes that the sizes of investor accounts are approximately equal. More accurately, we assume that of the total units of fund we are guaranteeing, over a year 25% of the guarantees are reset when it is optimal to do so.

[6]As mentioned above, a D.S.C. is applied to investors' accounts during the first five years of the contract. This effectively provides a penalty for lapsing during this time, as investors lose a percentage of the value of their accounts if they close them. This obviously reduces the chances of anti-selective lapsation.

4 Intuition and Heuristics

So far, we have claimed that the reset feature can make these contracts quite expensive to hedge. In Section 4.1 we provide some intuitive reasons as to why this is so. We also try to clarify some common misunderstandings regarding deferred payment schemes. In Section 4.2 we provide some heuristics for hedging positions which would be used in a dynamic hedging scheme.

4.1 Effect of the Reset Feature

Intuitively, a 10 year at-the-money put option provides a lower bound for the value of a segregated fund guarantee. This is because, if the reset feature is never used, we simply have a standard European put option.[7] Clearly, the segregated fund guarantee cannot be worth less than this 10 year option, but is it really worth more?

For simplicity, we begin our discussion by assuming that the cost of providing the guarantee is paid for up front. In other words, we consider the case of a single premium contract, with no deferred proportional fees. In Figure 1(a) we see that the cost of the providing a put option is monotonically decreasing in the asset value, whereas the cost of providing a segregated fund guarantee swings upwards for high asset levels. This happens because for high asset levels, it is optimal for the investor to lock in at a new higher guarantee level, and the value of the guarantee is relative to the current guarantee setting. As a result, the key difference when including the reset feature in a contract is that the contract should never be thought of as being out-of-the-money.

In order to fund the guarantee, it is typical to charge a proportional fee over the life of the contract. This approach attempts to overcome two problems which arise when the contract is paid for up front. First, the present value of the future incoming fees is also proportional to the current level of the underlying fund. This combats the effect described above of providing a new higher guarantee level. Second, unless the possibility that the guarantee level is reset to a higher level is taken care of appropriately in the hedging strategy, there is a potential for large hedging losses during the life of the contract. Consider the following simple example:

- An investor purchases a segregated fund. The underlying mutual fund is an index fund. The value of the index (and the guarantee level) is, say 10,000.
- To hedge downside risk, the insurer shorts the index (or does something similar, such as buying at-the-money put options or taking a short position in index futures contracts).
- Four days later, there is news that interest rates are likely to fall and the index rises abruptly to 10,300. At the same time, the investor resets his guarantee level.

[7] This ignores the mortality aspect of the contract. As noted earlier, for the investor demographic type we are considering, this feature does not have signficant value.

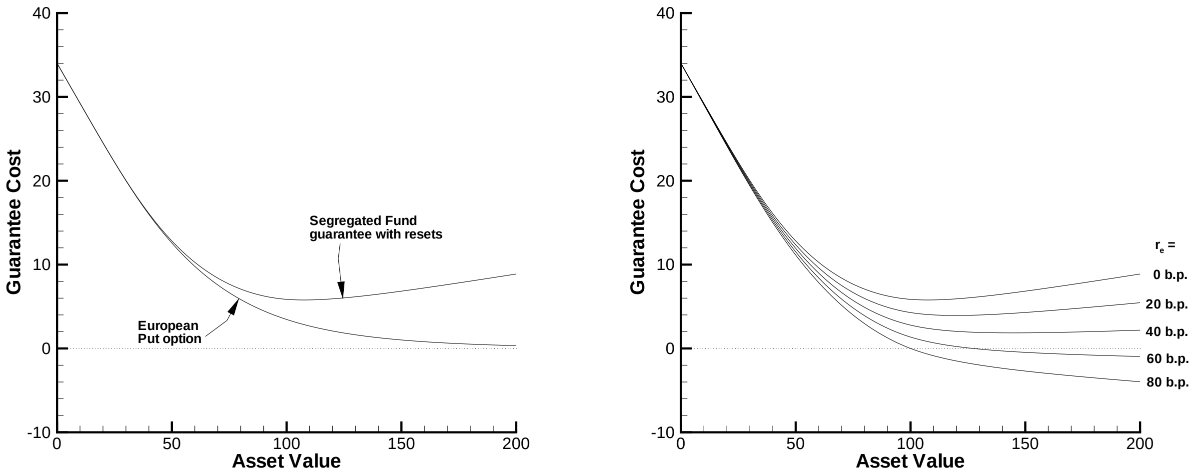

(a) *The effect of incorporating the reset feature in a segregated fund guarantee. The curve labeled as a segregated fund guarantee with resets provides two reset opportunities per year. The curve labeled as a European put option offers no resets. In these scenarios, no proportional fees are charged and the cost given is the initial value of the contract.*

(b) *The residual cost, net of future incoming fees, of providing the base scenario segregated fund guarantee when various proportional fees are charged. The correct proportional fee for this contract is $r_e = 80$ b.p., which results in the zero initial residual cost, net of incoming fees.*

FIGURE 1: *The cost of hedging a segregated fund guarantee with the reset feature.*

- The insurer then liquidates the original hedging strategy (which has lost money) and sets up a new hedge to cover the new guarantee level. However, the proportional fees collected from the investor over the four days will probably not cover the loss on the hedging position.

In other words, the insurer is exposed to a double-edged sword: the liability increases if the fund drops significantly (so that the initial guarantee is valuable), or if it rises significantly (so that the investor resets and receives a new at-the-money guarantee which is worth more than the original one). Clearly, it is more expensive to hedge this than a situation where the liability increases in only one direction. The idea behind charging a proportional fee over the life of the contract is that investors who utilize the reset feature will pay more in fees, since the maturity date of the contract is extended by 10 years upon reset.

So, instead of charging for the guarantee up front, we now receive a stream of incoming proportional fees. This changes our notion of the cost of providing the guarantee. This now refers to the required cost of setting up a hedge to cover the guarantee, *net of the future incoming proportional fees*. When no fees are provided to fund the guarantee, $r_e = 0$ b.p., the liability is simply the initial cost of the contract shown above (Figure 1(a)). In Figure 1(b) we show the required cost, net of incoming fees, of providing the guarantee as a function of the current asset level. As the proportional fee, r_e, is increased, the net residual cost of providing the guarantee decreases for higher asset levels. The relevant problem is to determine r_e so that at contract inception the net residual cost of providing the guarantee is zero. An important point to note is that as r_e is increased, the upward swing at high asset levels is pulled downwards and eventually has a profile which is similar to that of a put option. The implication is that using a deferred proportional fee *does* work, *provided* that the appropriate proportional fee is collected. However, concerns have been raised that some companies have sold these contracts for a fee which is too low.[8] In the next section we describe some of the consequences for the hedging scheme of undercharging using a deferred fee.

4.2 Hedging Heuristics

Suppose that we have a fixed M.E.R. and are interested in hedging the product, given the income stream from the deferred payments. We would like to develop heuristic approaches which allow us to qualitatively understand the required hedging for such a contract.

Here we will consider two cases for illustrative purposes: i) no fees are charged; and ii) the correct proportional fee (for the base scenario) of 80 b.p. is charged. Figure 2(a) plots the residual cost, net of future incoming fees, for these two cases. According to standard option pricing theory, given the value of an option, V, on an underlying asset with price S, one can establish a dynamic hedge by holding the fraction delta ($\Delta = \partial V/\partial S$) in the underlying asset and holding the remainder, $V - \Delta S$, in a risk free account. A graph of delta as a function of the current asset level is presented in Figure 2(b).

Note how the graph of the residual hedging cost behaves as the deferred fee which is collected varies (Figure 2(a)). When no fee is charged, the graph has three distinct regions.

- When the asset value is low, the cost of hedging is dominating by the hedging of the current guarantee level (which is in-the-money) and behaves like a put option. Here Δ will be negative and we should have a short position in a basket of stocks which simulates the underlying mutual fund.

- When the asset value is high, the curve turns upwards. Here, since Δ is positive, we hold a long position to protect ourselves against the investor locking in at a new higher level. Remember that in this scenario the option has been paid for up front and the hedging strategy must be able to support itself without the aid of incoming fees. We will extend this model to include fees below.

- In between (near at-the-money), we are not certain which of the two types of outcomes we should try and hedge. Here, $\Delta \approx 0$, and we put the hedging resources in a risk free account.[9]

[8] See Windcliff et al. (2001); Falloon (1999), among others.

[9] This assumes a "delta-neutral" hedging strategy. In theory, this is all that is needed, but it does assume continuous adjustment of the hedging position. In practice, positions can only be adjusted discretely. A more sophisticated approach which attempts to account for this is a "gamma-neutral" strategy. This would require holding a position in other options, even if the position has $\Delta \approx 0$. Readers interested in further details about this type of hedging should consult texts such as Hull (2000) or Wilmott (1998).

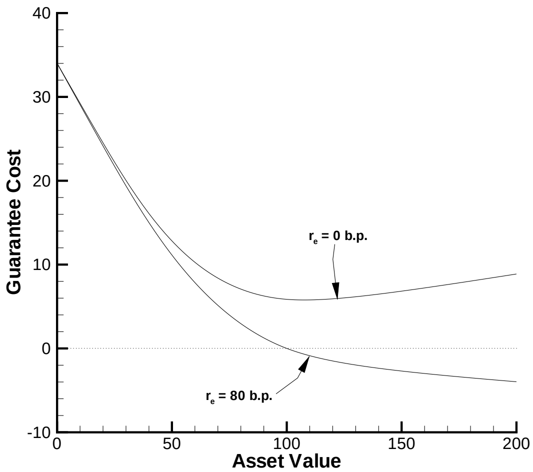

(a) *The residual cost, net of future incoming fees, of providing the base scenario segregated fund guarantee when there are no proportional fees charged, and when the correct proportional fee of $r_e = 80$ b.p. is charged.*

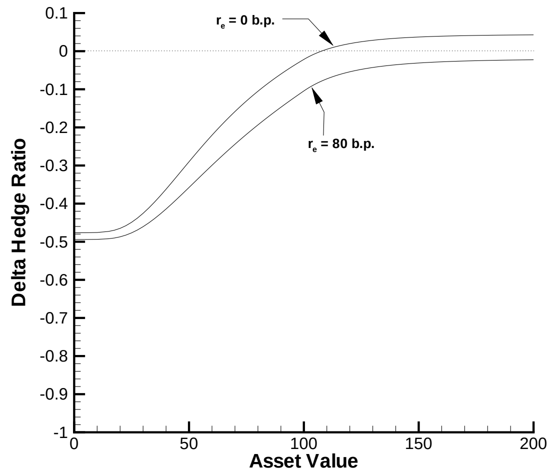

(b) *The delta hedging parameter, $\Delta = \frac{\partial V}{\partial S}$, for the base scenario segregated fund guarantee when no proportional fees are charged and when the correct proportional fee of $r_e = 80$ b.p. is charged.*

FIGURE 2: *The dynamic hedging of a segregated fund guarantee with the reset feature.*

> Notice that this is probably the most likely scenario (at least in a generally rising market) since at high asset levels the investor should reset and the contract would be at-the-money.

This qualitative behaviour can be seen in Figure 2(b) on the curve corresponding to $r_e = 0$ b.p.

As r_e is increased, the curve at higher asset levels is pulled downwards. Eventually, as a sufficiently large r_e is charged, the curve becomes monotonically decreasing. When this happens, the Δ hedging parameter (which corresponds to the slope of the curve of the value) is always negative, and hence we should always have a short position. We believe that most practitioners would consider this intuitive.

The problem arises from the fact that in some cases the deferred fee is not large enough to make the curve of the value monotonically decreasing. Here, we are somewhere between the two scenarios described above. In fact, it is possible that sometimes we will have to have a long position in the underlying when the asset level is high. This happens because the future incoming fees will not be sufficient to cover the hedging costs for the remainder of the contract. The important result to note is that we cannot ignore the reset feature when devising a hedging plan.

Another point needs to be made with regard to hedging. The graphs above of the hedging cost as a function of the asset value are snapshots taken at the time of the initial sale of the contract. Thus, the results given in this paper are only strictly applicable for some small time interval near the time of the initial sale date. In order to dynamically hedge these contracts, we really need the time evolution of the hedging costs as shown in Figure 3. Since this time evolution is very smooth, the general qualitative properties discussed above hold throughout much of the life of the contract.

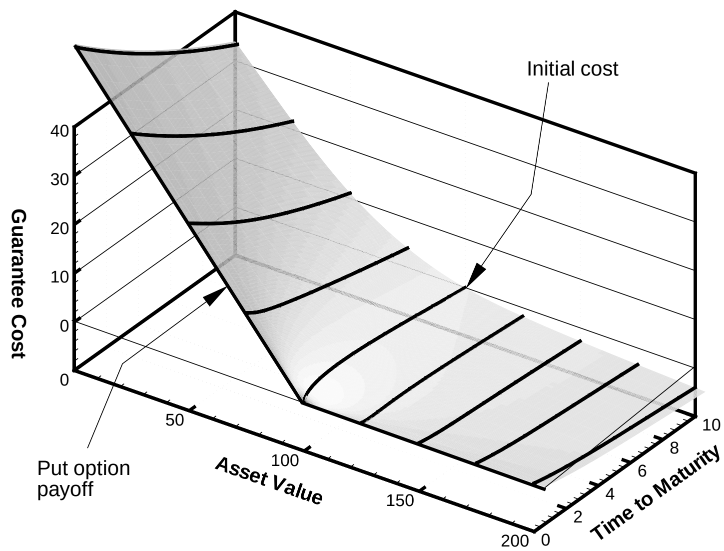

FIGURE 3: *The time evolution of the cost of providing the initial guarantee sold to the investor in the base case scenario where the correct proportional fee, $r_e = 80$ b.p. is being charged (note that the initial cost of providing the guarantee is zero). The time of sale of the contract corresponds to 10 years to maturity.*

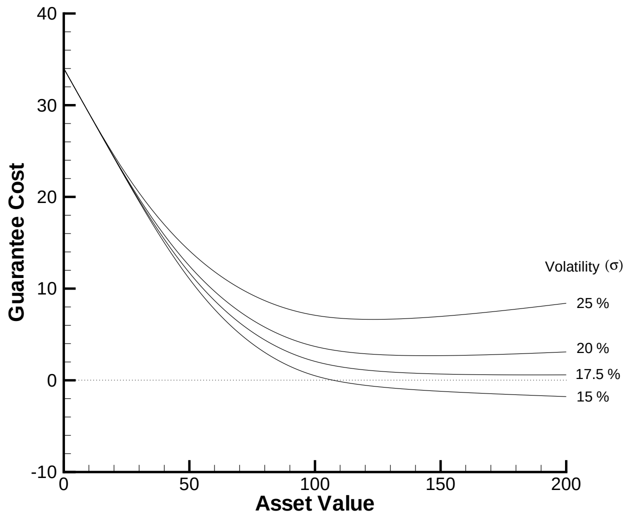

Volatility (σ)	Required Fee (r_e)
15%	55 b.p.
17.5%	80 b.p.
20%	105 b.p.
25%	155 b.p.

FIGURE 4: *The effect of volatility. Left: residual cost, net of future incoming fees, of providing a segregated fund guarantee when $r_e = 50$ b.p. Right: required fee to cover dynamic hedging costs.*

5 Variations

Our model takes several important input parameters from the market. For example, the cost of providing the guarantee will depend upon the volatility of the fund for which the guarantee is being provided. Other important factors include the risk free rate and the level of optimality displayed by investors.

5.1 Effect of Volatility and Risk Free Interest Rate

One very important parameter which we must estimate from the market when valuing these guarantees is the volatility of the underlying fund. In this paper, we give results for simple constant volatility models. The base scenario uses a volatility of $\sigma = 17.5\%$. Guarantees offered on more volatile funds can be dramatically underpriced. In Figure 4 we present a plot of the residual cost, net of future incoming fees, of providing a segregated fund guarantee given the fixed fee, $r_e = 50$ b.p. for various levels of volatility. In this figure we can see that for guarantees written on higher volatility funds, the residual cost, net of incoming fees, is positive. This indicates that the insurer will not be able to cover the costs of dynamic hedging with this choice of r_e. Figure 4 also contains a table which shows the proportional fee which should be charged in order to cover the hedging costs for the given volatilities.

Another important parameter in the model is the risk free interest rate. Figure 5 demonstrates that small changes in our assumptions about the risk free rate can have dramatic impacts on the cost of providing these guarantees. It is worth reiterating here that our assumptions that the interest

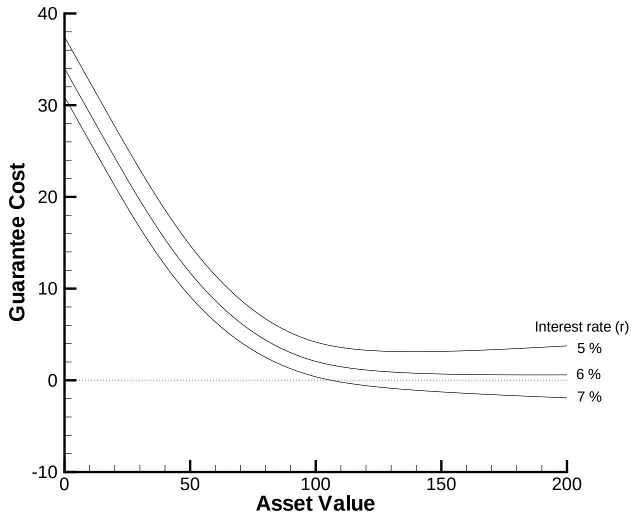

Interest Rate (r)	Required Fee (r_e)
5%	115 b.p.
6%	80 b.p.
7%	55 b.p.

FIGURE 5: *The effect of the risk free interest rate. Left: residual cost, net of future incoming fees, of providing a segregated fund guarantee when $r_e = 50$ b.p. Right: required fee to cover dynamic hedging costs.*

rate and volatility are constant are clearly unrealistic, especially given the long term nature of these contracts. In practice, we would have to carefully consider how to estimate these parameters in order to minimize the effects of these simplifying assumptions. For example, we would probably want to use a volatility parameter that is consistent with prevailing implied volatilities from options markets, rather than the historical volatility of the underlying funds in question.

We should point out that these results might be viewed as conservative in that we are assuming a degree of investor optimality of 25%. We discuss the effect of investor behaviour in the next section. Typically, segregated funds are sold with a value of r_e in the range of 40 to 80 b.p. In view of the results in Figures 4 and 5, we can see that these fees are appropriate for low to moderate volatility funds and interest rates that are not very low, even when we assume that investors do not use their reset options with a high level of efficiency.

5.2 Effect of Investors' Behaviour

Another very important parameter which our model uses is the level of optimality displayed by investors when using the reset feature. Strictly speaking, from a no-arbitrage viewpoint, one *should* assume that investors act optimally. In practice, most investors do not act optimally and we would like to incorporate this into our model.

One should note that by assuming that investors act non-optimally and charging fees based on this, insurers expose themselves to three additional sources of risk. First, these contracts are extremely long term and it is possible that there may be secular changes in investors' behavior over

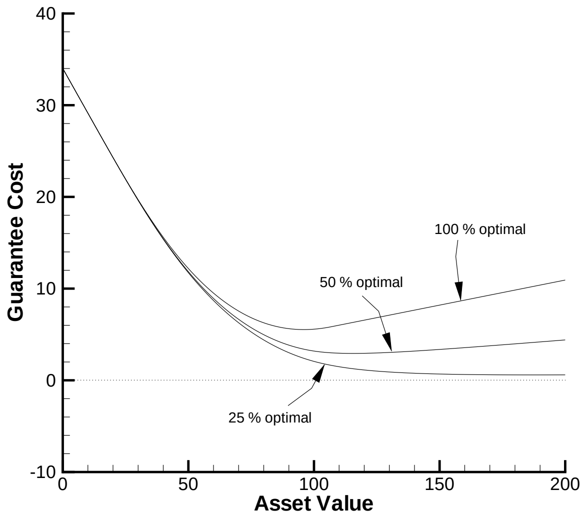

Investor Optimality	Required Fee (r_e)
25%	80 b.p.
50%	95 b.p.
100%	130 b.p.

FIGURE 6: *The effect of investor optimality. Left: residual cost, net of future incoming fees, of providing a segregated fund guarantee when $r_e = 50$ b.p. Right: required fee to cover dynamic hedging costs.*

the life of the contract. Similar changes are known to have occurred in other contexts, for example prepayment rates on mortgages in the U.S. Second, there may be considerable uncertainty about the current level of investor optimality, since few insurers appear to be in a position to measure this with any degree of rigor. Third, it may be possible for a single knowledgeable investor to take a relatively large position, affecting our assumption about the level of optimality displayed by investors.

As shown in Figure 6, the degree of optimality displayed by investors has a large impact on the required hedging costs of these contracts. Why is it so expensive to hedge a contract which is reset optimally? Investors who reset optimally are more likely to catch the peaks in the market, and thus they end up with higher guarantee levels. Alternatively, they are likely to (anti-selectively) lapse much sooner, thus depriving the insurer of fee income.

5.2.1 Optimal Exercise Boundary

Collecting data to determine quantitatively the level of optimality is difficult and requires knowledge of the optimal exercise boundary. In other words, we could compare data on investors' use of their reset options with the optimal exercise boundary to assess whether the appropriate optimality level should be 25% (as assumed in our base scenario above) or some other number. At this stage, we do not have sufficient data on the use of the reset provision by investors to do this. However, we are able to compute the optimal exercise boundary. Figure 7 illustrates this for a contract which allows the investor one reset per annum. Note that this boundary applies only to the initial contract sold

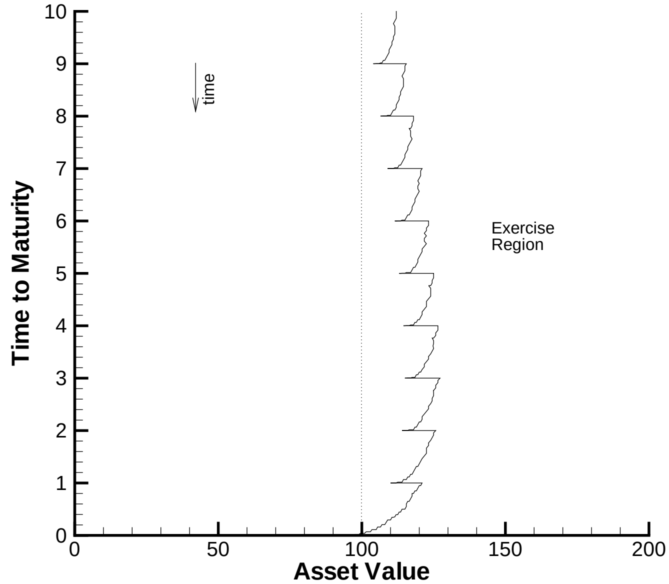

FIGURE 7: *The optimal exercise boundary of a segregated fund guarantee which allows one reset per annum, assuming a proportional fee of $r_e = 50$ b.p. Note that this is for the initial contract sold, where the investor has not reset. Once the investor does reset, the exercise boundary changes.*

to the investor, i.e. for the first use of a reset. Once the investor resets, the exercise boundary changes. A rough heuristic for the contract specifications and market parameters used suggests that these contracts should be reset when the value of the fund is approximately 10% to 25% above the current guarantee level.[10]

The location of the exercise boundary depends on the current maturity date of the contract. We can see that there is a trade-off between getting a higher guarantee level by resetting and deferring the maturity date of the contract by another ten years. The jumps in the location of the optimal exercise boundary occur because the investor receives a new reset opportunity each year.

5.2.2 Effect of Deterministic Lapsing

As mentioned above, for a variety of reasons investors may withdraw from these contracts independently of the value of the underlying fund. In Figure 8 we compare the net value of these contracts for different levels of investor deterministic lapsing. As is to be expected, an increase in the deterministic lapse rate reduces the net cost of hedging these contracts because fewer people are remaining in the fund to collect any payoffs made by the guarantee. However, as seen in

[10]In a report prepared by a Canadian Institute of Actuaries task force (Canadian Institute of Actuaries, 2000), it is suggested that contracts can be viewed as being "optimally" reset when the value of the underlying fund is more than 15% above the guaranteed level.

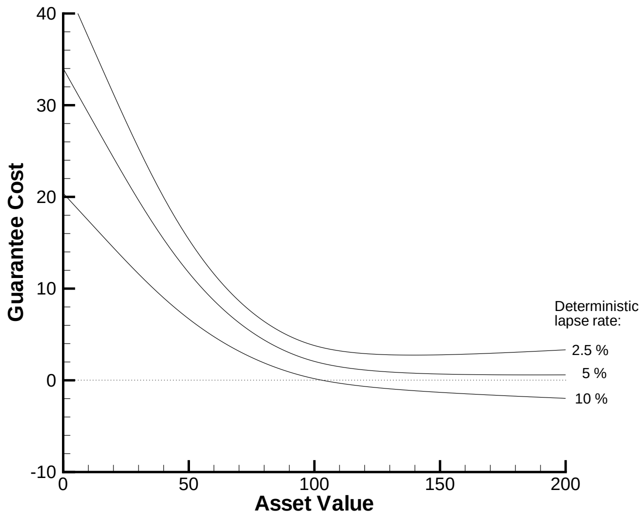

Deterministic lapse rate	Required Fee (r_e)
2.5%	100 b.p.
5%	80 b.p.
10%	55 b.p.

FIGURE 8: *The effect of deterministic lapsation. Left: residual cost, net of future incoming fees, of providing a segregated fund guarantee when $r_e = 50$ b.p. Right: required fee to cover dynamic hedging costs.*

the accompanying table, the required proportional fee does not decrease as rapidly as might have been anticipated. This is because, as the rate of deterministic lapsing is increased, the number of investors remaining in the fund (and thus the value of the future incoming fees) decreases.

Note that the lapsation model used is very simple; some given fraction of the investors lapse out of the contract every year. This approach probably undervalues these contracts when the guarantee is deep-in-the-money. In this situation, investors would be better off remaining in the segregated fund and receiving their minimum locked in value rather than lapsing and throwing away their guarantee.

5.2.3 Anti-Selective Lapsation

Because of investors' ability to lapse anti-selectively, institutions that have removed the reset feature from their product offerings may not have reduced their risk exposure as much as one might initially believe. Consider a case where there are no reset provisions or mortality benefits, so we have just a ten year European put option. Assume for simplicity that the correct proportional fee is being charged for this contract, and that there are no back end fees, so that investors can lapse without penalty. In the absence of explicit reset options, investors have the ability to create resets synthetically by lapsing and then re-entering the contract. Why then does it appear that the explicit reset option creates so much additional value? The reason is that we haven't adequately valued the ability of investors to perform resets synthetically. When investors are not charged a single up front fee, they have a valuable type of default option: they can avoid continuing to pay for a deep out-of-the-money put and acquire a new at-the-money guarantee with no additional expense. If we simply think of these contracts as a ten year puts and ignore this default option, we significantly undervalue them. Our model which incorporates resets is effectively including this option, and this is why it appears to generate much higher values.[11] Of course, in practice various factors do constrain investors' ability to follow this synthetic reset strategy. Back end fees are levied in the first few years. There may be tax consequences associated with withdrawing funds from an account. Moreover, to the extent that the proportional fees being charged are too low, investors have an incentive to remain in their existing contracts. However, our main point is that removing explicit reset provisions may not be a panacea if investors can effectively reset on their own.

6 Contract Modifications

To this point, we have seen that it can be very expensive to hedge contracts which include the reset feature. In this section, we look at possible approaches which may reduce the cost of providing these guarantees. Specifically, we explore some ways of modifying these contracts so that they remain attractive to the investor, yet are less expensive to hedge in the market.

At the outset, we should observe that in our research we have investigated many possible modifications of these contracts. It has been our experience that limiting contract features which depend on the optimality of investors' actions does not reduce the value of these contracts as much as might be expected.

6.1 Effect of Number of Reset Opportunities

One example of a contract modification which depends on the optimality displayed by investors is altering the number of reset opportunities per annum. Clearly, the value of the contract must decrease as the number of reset opportunities decreases. In Figure 9 we show the cost of providing a segregated fund guarantee with one to four reset opportunities per annum, for both 25% and 100% levels of investor optimality.

If we look at the lower collection of curves, corresponding to a 25% optimal investor profile, the cost does decrease as expected, but the curves are virtually indistinguishable. Further, in the accompanying table we see that the required proportional fee is virtually unaffected and is approximately 80 b.p. for either one, two or four resets allowed per annum. Why is the effect of this contract modification so small? The reason is that we are already assuming a high degree of investor non-optimality. In effect, we are limiting a feature which we are assuming the investors are not using efficiently anyways.

[11] In fact, the synthetic reset strategy could be executed at any time. Therefore, the number of possible resets is unlimited (unlike in our model).

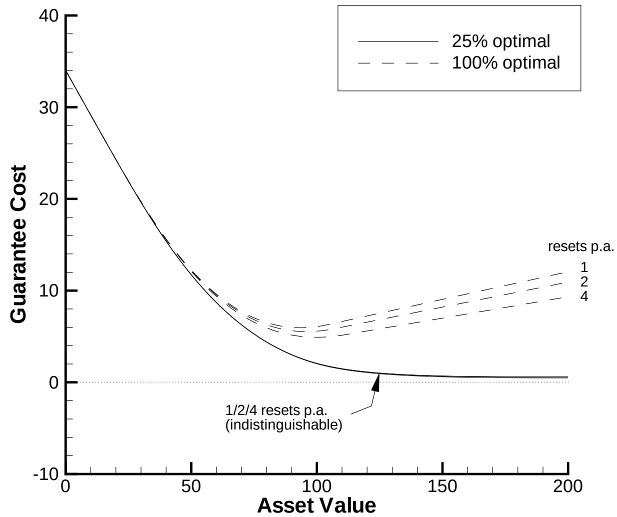

Number of Resets (p.a.)	Required Fee (r_e)
25% optimal behaviour	
1	80 b.p.
2	80 b.p.
4	80 b.p.
100% optimal behaviour	
1	120 b.p.
2	130 b.p.
4	135 b.p.

FIGURE 9: *The effect of the number of resets available per annum. Left: residual cost, net of future incoming fees, of providing a segregated fund guarantee offering one, two and four reset opportunities per annum, when $r_e = 50$ b.p., assuming 25% and 100% investor optimality. Right: required fee to cover dynamic hedging costs.*

In the upper collection of curves in Figure 9, corresponding to 100% optimal investor behaviour, the differences become noticeable. This is also reflected in the accompanying table. Thus, reducing the number of reset opportunities slightly reduces the risk exposure to investors who make effective use of their optionality while maintaining a product with the attractive ability to lock in market gains for risk averse investors.

6.2 75% Guarantee Level

One approach which has been implemented by many institutions is the offering of a 75% guarantee level. Here the investor is provided with a guarantee at 75% of the original principal invested and upon reset the guarantee level is set to be 75% of the asset value at the time of the reset. Offering this reduced guarantee level does make these contracts much less expensive to hedge. In Figure 10 we compare the cost of establishing a hedging strategy, net of future incoming fees, for a contract which offers a reduced 75% guarantee with the full guarantee offered in the base case. Notice that for a proportional fee of 50 b.p. the full guarantee is a liability under the market conditions $\sigma = 17.5\%$, $r = 6\%$ assuming a 25% optimal investor profile. On the other hand, under these same conditions, the 75% guarantee can be offered at a profit by the insurer. The proportional fee which is required to cover the insurance costs is 25 b.p., as seen in the accompanying table.

Since this contract variation is quite common, we explore the effects of variations in volatility, interest rate and investor optimality for 75% guarantees. In Table 2 we observe that the reduction in the cost of providing the 75% guarantee decreases for higher volatility funds when compared

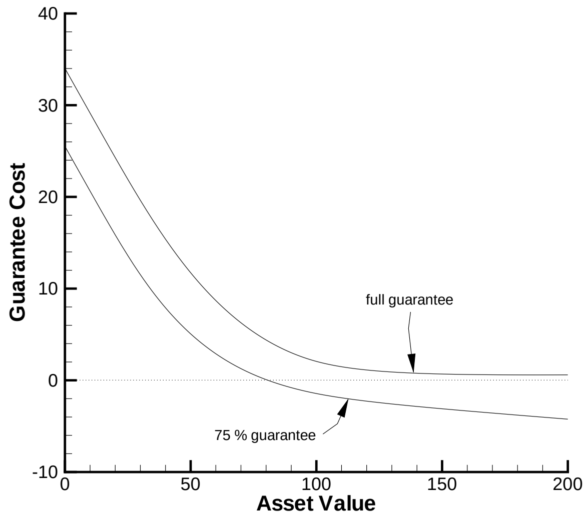

Contract Type	Required Fee (r_e)
Full guarantee	80 b.p.
75% guarantee	25 b.p.

FIGURE 10: *The effect of the guarantee level. Left: residual cost, net of future incoming fees, of providing a segregated fund guarantee when $r_e = 50$ b.p., $\sigma = 25\%$, and $r = 6\%$. Right: required fee to cover dynamic hedging costs.*

with the results for the full guarantee. For a low volatility fund where $\sigma = 15\%$, the required proportional fee is smaller by a factor of four. However, for a high volatility fund with $\sigma = 25\%$, the required proportional fee is only smaller by a factor of two. This occurs because a more volatile fund has a greater likelihood of large swings in asset value which are required in order for the 75% guarantee to be in-the-money at maturity. The effect of interest rates on the cost of providing a 75% guarantee shown in Table 2 is similar that found for the full guarantee contract.

Table 2 also illustrates that offering a 75% guarantee reduces the value of the contract's sensitivity to the degree of optimality shown by investors when compared with the results for the full guarantee in Figure 6. This occurs because the contract received upon reset has been decreased in value by the restricted 75% guarantee. Since the reset feature will be used less for a 75% guarantee, the level of optimality displayed by investors plays less of a role in valuing these contracts.

7 An Efficient Approximation for Valuing Segregated Fund Guarantees

In order to rigorously model the reset feature contained in segregated funds, it is necessary to maintain state variables corresponding to the number of reset opportunities used during the current year, and the current maturity date of the contract in addition to the current value of the underlying mutual fund. Moreover, although it is unnecessary in the geometric Brownian motion setting we are using here, in some circumstances an additional state variable may be required for the current

Volatility (σ) ($r = 6\%, \alpha = 25\%$)	75% guarantee	Full guarantee (from Figure 4)
15%	15 b.p.	55 b.p.
17.5%	25 b.p.	80 b.p.
20%	40 b.p.	105 b.p.
25%	70 b.p.	155 b.p.
Interest rate r ($\sigma = 17.5\%, \alpha = 25\%$)	75% guarantee	Full guarantee (from Figure 5)
5	35 b.p.	115 b.p.
6	25 b.p.	80 b.p.
7	15 b.p.	55 b.p.
Investor optimality ($\sigma = 17.5\%, r = 6\%$)	75% guarantee	Full guarantee (from Figure 6)
25%	25 b.p.	80 b.p.
50%	30 b.p.	95 b.p.
100%	45 b.p.	130 b.p.

Table 2: *Effect of volatility, interest rate and investor optimality on the proportional fee required for a 75% maturity guarantee.*

guarantee level (see Windcliff et al., 1999, for further discussion). This results in a three (or more) dimensional system which is time-dependent. Solving this numerical problem can be quite expensive computationally. However, due to the structure of the mathematical problem very efficient parallel algorithms can be implemented to reduce the valuation time on multiprocessor computers (see Windcliff et al., 2000).

An alternative approach which reduces the dimensionality of the numerical problem (and hence the computation time) is to build on the intuition described in section 5.2.3. As we have seen, a contract that allows anti-selective lapses but no resets may be almost as valuable as a contract that does allow resets. It follows that we can derive an approximate value for the latter contract by modelling the former. Instead of explicitly modelling the reset feature, we simply model anti-selective lapses. In effect we assume that the contract is correctly priced, and that the value of the contract (net of future fees) is zero at the point of reset. In other words, we assume that it makes no difference to the insurer whether the investor resets or lapses. If the value of the contract were to fall below zero, then the investor should reset, thereby obtaining a contract with zero net worth.

This approximation can be very accurate, as shown in Figure 11(a) for our base scenario. Further, the complexity of the numerical computation is equivalent to pricing a (long term) standard American put option. This can readily be handled using lattice methods such as binomial or trinomial trees or finite difference methods. These techniques allow for very rapid computation and are also quite straightforward to implement.

However, one must be careful when making such approximations. If the correct proportional fee is not being charged, the net value of the ten year guarantee is not zero, contradicting the assumption made in the approximation. In this case, the heuristic handling of the reset feature can

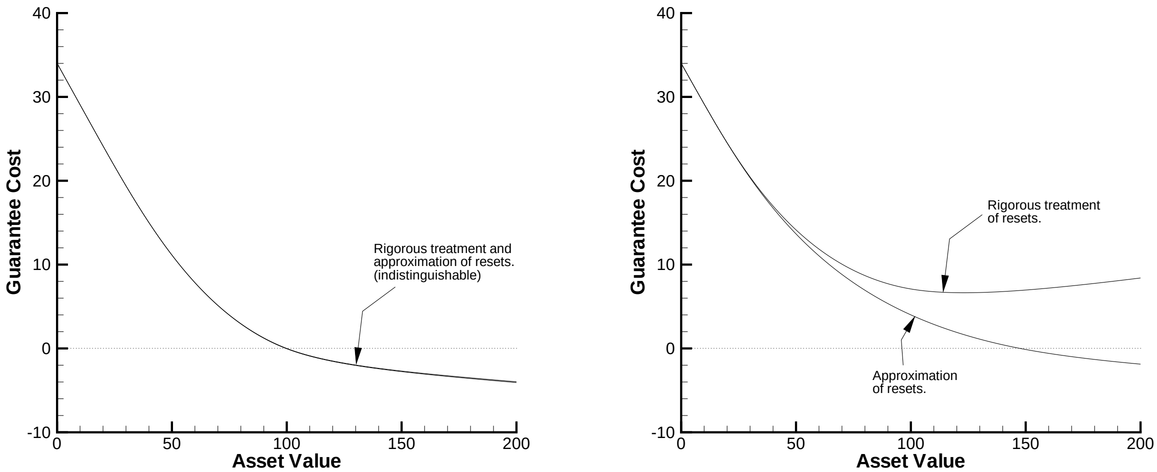

(a) *Approximation of the reset feature when the correct proportional fee is charged. Base case scenario with $r_e = 80$ b.p.*

(b) *Approximation of the reset feature when an incorrect proportional fee is charged. Base case scenario with $\sigma = 25\%$ and $r_e = 50$ b.p.*

FIGURE 11: *Results computed using an approximation of the reset feature.*

severely underestimate the value. One example is shown in Figure 11(b), where we attempt to price a guarantee on a high volatility fund, $\sigma = 25\%$, using a proportional fee of $r_e = 50$ b.p.[12] Note also that in situations like this our estimates of hedging parameters can also be very inaccurate.

8 Conclusions

In this paper we discuss some of the financial implications of pricing and hedging segregated fund guarantees which include the reset feature. Some of the key results include:

- Segregated fund guarantee products which offer the reset feature appear in some cases to be underpriced by Canadian insurers when valued using a no-arbitrage approach.
- We have presented intuitive reasoning regarding why it can be quite expensive to hedge the reset feature, and presented some heuristics for establishing a hedging strategy.
- We have also presented rigorous numerical results, based on the model developed in Windcliff et al. (2001), which show appropriate proportional fees as well as the cost of providing a segregated fund guarantee (net of future incoming fees) for a given fixed fee structure.
- We have investigated the possibility of modifying these contracts so that they remain attractive to investors yet are less expensive to hedge. We note that these contracts do not become

[12] The correct proportional fee, from Figure 4, is $r_e = 155$ b.p.

less expensive to hedge by limiting contract features which we are assuming are being used non-optimally.

- Limiting the value of the guarantee to 75% of the current level of the asset does reduce the hedging cost substantially. The value of the resulting contract is relatively less sensitive to optimal investor behaviour, but more sensitive to high volatility.

- If the correct proportional fee, r_e, is being charged, handling resets by imposing a minimum value constraint at zero can be an accurate approximation. This approximation allows these contracts to be valued in the time required to value a standard American put option. However, care must be taken if the proportional fee being charged is too small.

The results given in this paper can be viewed as the minimum cost of providing these guarantees since we are modeling a contract which has much less optionality than many existing marketed contracts and we have assumed a low degree of investor optimality. Also, additional reserves may be required to handle unhedged risks such as basis risk, volatility risk, and interest rate risk as well as transactions costs.

References

Boyle, P. P. and M. R. Hardy (1997). Reserving for maturity guarantees: Two approaches. *Insurance: Mathematics & Economics* 21, 113–127.

Boyle, P. P., A. W. Kolkiewicz, and K. S. Tan (1999). Valuation of the reset option in segregated fund contracts using quasi-Monte Carlo methods. University of Waterloo Institute of Insurance and Pension Research Report 99-10.

Boyle, P. P. and E. S. Schwartz (1977). Equilibrium prices of guarantees under equity-linked contracts. *Journal of Risk and Insurance* 44, 639–660.

Brennan, M. J. and E. S. Schwartz (1976). The pricing of equity-linked life insurance policies with an asset value guarantee. *Journal of Financial Economics* 3, 195–213.

Canadian Institute of Actuaries (2000). Report of the CIA Task Force on Segregated Fund Investment Guarantees. `http://www.actuaries.ca/publications/2000/20020e.pdf`.

Cox, J. C. (1975). Notes on option pricing I: Constant elasticity of variance diffusions. Working paper, Stanford University (reprinted in *Journal of Portfolio Management* 22 (1996), 15–17).

Falloon, W. (1999, August). Canada's option nightmare. *Risk* 12, 60.

Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies* 6, 327–343.

Hull, J. C. (2000). *Options, Futures, & Other Derivatives* (4th ed.). Prentice-Hall, Upper Saddle River, NJ.

Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics* 3, 125–144.

Rabinovitch, R. (1989). Pricing stock and bond options when the default-free rate is stochastic. *Journal of Financial and Quantitative Analysis* 24, 447–457.

Wilmott, P. (1998). *Derivatives: The Theory and Practice of Financial Engineering.* John Wiley & Sons, West Sussex, England.

Windcliff, H., P. A. Forsyth, and K. R. Vetzal (1999). Shout options: A framework for pricing contracts which can be modified by the investor. *Journal of Computational and Applied Mathematics*, forthcoming, `http://www.scicom.uwaterloo.ca/~paforsyt/shoutnum.ps`.

Windcliff, H., P. A. Forsyth, and K. R. Vetzal (2001). Valuation of segregated funds: Shout options with maturity extensions. *Insurance: Mathematics & Economics*, forthcoming, `http://www.scicom.uwaterloo.ca/~paforsyt/seg.ps`.

Windcliff, H., K. R. Vetzal, P. A. Forsyth, A. Verma, and T. F. Coleman (2000). An object-oriented framework for valuing shout options on high-performance computer architectures. *Journal of Economic Dynamics and Control*, forthcoming, `http://www.scicom.uwaterloo.ca/~paforsyt/parallel.ps`.