Hedging Segregated Fund Guarantees

Heath A. Windcliff
Dept. of Computer Science
University of Waterloo, Waterloo ON, Canada N2L 3G1.
hawindcliff@elora.math.uwaterloo.ca

Peter A. Forsyth
Dept. of Computer Science
University of Waterloo, Waterloo ON, Canada N2L 3G1.
paforsyt@elora.math.uwaterloo.ca

and

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

February 18, 2002.

Abstract

Segregated funds have become a very popular investment instrument in Canada. Segregated funds are essentially mutual funds that have been augmented with additional insurance features, which provide a guarantee on the initial principal invested after a specified time horizon. They are similar in many respects to variable annuities in the United States. However, segregated funds often have complex embedded optionality. For instance, many contracts provide a reset provision. This allows investors to increase their guarantee level as the value of the underlying mutual fund goes up. These contracts typically also offer features such as mortality benefits, where the guarantee is paid off immediately upon death of the investor. These contracts are further complicated by the possibility of investor lapsing since the payment for the guarantee is amortized over the life of the contract. In this work we describe hedging strategies which allow underwriting companies to reduce their risk exposure to these contracts. The hedging techniques incorporate the strengths of both actuarial and financial approaches. In particular, we look at some of the difficulties that arise due to the fact that in many cases the underwriting company is not able to take short positions in the underlying mutual fund. An alternative is to hedge using other actively traded securities, such as index participation units, stock index options, or stock index futures contracts. However, due to the mismatch between the hedging instrument and the segregated fund contract being hedged, there is additional basis risk. We investigate the performance of these types of hedging strategies using stochastic simulation techniques and investigate the return on the initial capital outlay mandated by new regulatory requirements.

Segregated funds have become a very popular investment instrument in Canada. Segregated funds are essentially mutual funds that have been augmented with additional insurance features that provide a guarantee on the initial principal invested after a specified time horizon. They are similar in many respects to variable annuities in the United States. However, segregated funds often have complex embedded optionality. For instance, many contracts provide a reset provision. This allows investors to increase their guarantee level as the value of the underlying mutual fund goes up. These contracts typically also offer features such as mortality benefits, where the guarantee is paid off immediately upon the death of the investor. Furthermore, because the payment for the guarantee is usually amortized over the life of the contract, there are additional complications due to investor lapsing.

In Canada, the Office of the Superintendent of Financial Institutions (OSFI) has recently imposed stricter new regulations for these contracts, requiring that insurers set aside a substantial amount of capital to back up these guarantees. These reserve requirements can be reduced if appropriate hedging strategies have been put in place. In this work we describe hedging strategies which allow underwriting companies to reduce their risk exposure to these contracts. The hedging techniques incorporate the strengths of both actuarial and financial approaches. We investigate the performance of these types of hedging strategies using stochastic simulation techniques.

It is well known that a contract with a maturity guarantee attached to it can be thought of as an investment in the underlying asset combined with a put option, which may have very complex features such as mortality benefits or reset provisions. Hedging strategies that reduce downside risk for put options typically involve shorting the underlying mutual fund. For obvious regulatory reasons, it may not be possible for the underwriting company to take short positions in the underlying mutual fund. One alternative is to set aside a bond with face value equal to the guarantee level and dynamically hedge a (possibly complex) call option position. The advantages of this formulation are that the insurer can easily take long positions in the underlying mutual fund and downside risk has been completely hedged. However, now the insurer is exposed to a considerable amount of upside risk if the hedging of the call option does not work as planned. Another alternative is to hedge using other actively traded securities, such as index participation units, stock index options, or stock index futures contracts. However, due to the mismatch between the hedging instrument and the segregated fund contract being hedged, there is additional basis risk. Using stochastic simulation we can quantify the risk associated with selling these contracts in a more realistic setting which includes non-optimal investor behaviour and basis risk.

We would like to emphasize that the decision of whether or not to hedge these contracts in many situations is a management issue. In some cases, hedging may be necessary to reduce the capital requirements due to the new regulations. The hedging strategies described here are capable of reducing the downside risk associated with writing these contracts. Of course this comes at an expense, as the expected profit of the hedged position is lower than the expected profit of the unhedged position. The purpose of this paper is to develop the tools required to make an informed decision as to whether or not an institution would benefit from implementing a hedging strategy for these contracts.

Description of the Segregated Fund Contract

The term, segregated fund, often refers to a mutual fund combined with a long-term maturity guarantee (typically 10 years) with additional complex features. One popular provision that is included with many of these contracts is the reset feature. When the investor resets, they exchange their existing guarantee for a new 10-year maturity guarantee set at the current value of the mutual fund. Hence, the reset feature allows the investor to lock in market gains as the value of the underlying mutual fund increases. Typically, the investor is able to reset the contract up to a maximum of two or four times per calendar year. This introduces an optimization component to these contracts where the investor must decide when they should reset and lock in at the higher guarantee level.

In addition to the reset feature, many other exotic features are included in segregated fund guarantees. For example, many segregated funds include a death benefit so that the guarantee is paid out immediately, even if the investor expires before the maturity date of the guarantee. As the investor ages, the mortality benefits may become more valuable and typically resets are not permitted after, say, the investor's 70th birthday. Alternatively, more complex variations of the reset feature can also be introduced as the investor becomes older. For example, after the investor's 70th birthday the guarantee level upon reset may be some fraction of the value of the underlying mutual fund at the time of the reset. In practice the investor is not charged an initial fee for the segregated fund guarantee. Instead the investor pays a higher management expense ratio (MER) over the life of the contract to cover the cost of providing the guarantee. The total MER can be considered to be the sum of the proportional fee, r_m, which is allocated to the management of the underlying mutual fund, together with the proportional fee, r_e, which is allocated to fund the guarantee portion of these contracts.

It may be optimal for the investor to lapse and avoid paying the higher MER if the guarantee is unlikely to be in-the-money at maturity. For these and other reasons, such as need for liquid assets, an investor may withdraw their investment from the segregated fund contract.

The reset feature described above will help to reduce the amount of investor lapsing since the guarantee can be reset to a new at-the-money guarantee. Further, a proportional deferred sales charge (DSC) is often applied if the investor withdraws the investment during the first several years. Based on communications with vendors of these contracts, it appears as though this fee is paid to the underlying mutual fund and that in practice none of this fee is allocated to fund the guarantee portion of the contract. It should be reiterated that investor lapsing is not always beneficial to the insurance company writing the guarantee portion of these contracts. Specifically, since the payment of the guarantee is deferred over the life of the contract, any hedging costs incurred by the insurance company may not be recovered if the investor lapses prematurely.

The numerical experiments presented in this paper will be based on two contracts that are described in Table 1. The first contract is a simple 10-year maturity guarantee with no reset provisions. The second contract is a prototypical segregated fund guarantee that incorporates the reset feature. Both contracts provide mortality benefits so that the guarantee is paid out immediately upon the death of the investor. No initial fee is charged to enter into these contracts and the investor pays for these guarantees by the increased MER as described in the table with the sliding scale DSC to mitigate investor lapsing. Table 1 also describes the key market parameters used in the simulations such as the risk-free interest rate and volatility of the underlying mutual fund.

Table 1 approximately here.

Due to the complexity of these contracts it is difficult to draw general conclusions from individual numerical experiments. The numerical results in this paper are intended to provide a study of the behaviour of a realistic contract but small changes in the contractual details or variations in the market settings such as the volatility and risk-free interest rate will affect the pricing and hedging of these contracts. Interested readers are referred to (Windcliff et al. 2002) for the effects of these and other parameters, such as the level of investor optimality, on the valuation of segregated fund guarantees.

The Distribution of Returns for Unhedged Positions

In order to quantify the risk involved with writing the segregated fund guarantees described in the previous section, we can investigate the distribution of returns for an unhedged position. At this point, we will not assume optimal investor behaviour, but will presume that investors use heuristic rules for the reset feature and optimal lapsing. We will find that there can be a substantial amount of downside risk to the insurer when writing segregated fund guarantees with a reset provision, even when investors act non-optimally.

Let S represent the value of the underlying mutual fund and let K be the current guarantee level. In this section we will use the following heuristic rules for applying the reset feature and lapsing.

- Heuristic reset rule: Investors will reset the guarantee level if there are reset opportunities remaining and $S > 1.15K$; i.e. if the value of the underlying mutual fund has risen so that the current guarantee level is 15% out-of-the-money.
- Heuristic lapsing rule: Investors will lapse out of the contract at time t^*, and thereby avoid paying the remaining proportional fees, if there are no remaining reset opportunities at time t^* and $S > 1.4K$; I.e. if the value of the underlying mutual fund has risen so that the guarantee level is 40% out-of-the-money.

Investigation of the optimal reset region shows that resetting the guarantee when the underlying asset has risen by 15% can be a reasonable approximation to the optimal exercise boundary during the first few years of the contact or during the first few years after a reset has taken place (see (Windcliff et al. 2002)). In fact, this heuristic rule has been adopted by a Canadian Institute of Actuaries task force on segregated funds during their assessment of risk management strategies.

The heuristic lapsation rule is based on a plot of the optimal lapsation boundary given in (Windcliff et al. 2001). When the investor has no remaining reset opportunities they may be better off lapsing out of the contract to avoid paying the proportional fees since the guarantee is unlikely to be in-the-money at maturity. In fact, as discussed in (Windcliff et al. 2001), even if reset features are not explicitly offered, investors can synthetically create them by lapsing and re-entering the contract, thereby obtaining a new at-the-money guarantee.

In the Monte Carlo simulations provided in this paper, it is assumed that the investor makes decisions regarding the reset feature and optimal lapsing 100 times per year, or approximately twice every week. Numerical experiments indicate that more frequent exercise decisions by the investor do not appreciably affect the results.

To quantify the risk associated with writing an unhedged segregated fund guarantee, we consider the 95% conditional tail expectation (CTE) and the annualized rate of return on an initial capital requirement. Recent regulatory changes from OSFI have introduced stricter capital requirements for companies offering these contracts to ensure that sufficient resources are available to back up these guarantees. Specifically, if no hedging strategy is put in place OSFI requires that the insurer set aside the 95% CTE in liquid, risk-free instruments. The 95% CTE is the expected value of the outcomes that lie in the worst case 5% tail. In other words, the 95%

CTE is the mean value of the worst case outcomes that are ignored by a 95% value at risk (VaR) measurement. In comparison with VaR measures, the CTE is much more conservative when setting aside capital for contracts that exhibit a long tail of values, which occur with relatively low probability, such as segregated fund guarantees. If the insurer implements a hedging strategy, the OSFI capital requirement can be reduced by up to a maximum of 50% of the reduction in the 95% CTE indicated by the proposed hedging strategy. The capital must be invested in safe, liquid instruments, and for this paper we will assume that the capital investment grows at the risk-free rate. It should be pointed out that we have chosen the proportional fee r_e so that the cost of hedging, net of future incoming fees for these contracts is initially zero; in other words, the reserve amount is zero. As a result the total balance sheet requirement and capital requirements are identical.

We estimate the required capital by generating simulations of the mutual fund path, thereby generating a profit and loss (P&L) distribution for the writer of the guarantee as shown in Figure 1. The P&L for a simulation is given by

$$\text{P \& L} = (\text{hedge value} - \text{payoff value})\big|_{t^*} \times e^{-\bar{r}t^*}$$

where $\bar{r}$ is the discounting rate used and t^* is the time that the contract is terminated. If we use the risk-free rate as the discounting rate, $\bar{r} = r$, then the P&L distribution can be used to estimate the 95% CTE. This is the amount that must be set aside in risk-free instruments so that the insurer has sufficient resources to back up the guarantee in the average of the worst case situations.

Figure 1 approximately here.

Table 2 gives statistics for the P&L distribution. It is difficult to draw useful comparisons between P&L distributions since there are many risk/reward tradeoffs to consider and the duration of the contract is uncertain due to investor lapsing, mortality and the investor's use of the reset feature. In this work we use the distribution of the annualized return on the capital set aside for these guarantees as a measure of the profitability of offering these products. The outlay of capital to satisfy the OSFI requirements can be thought of as introducing an associated cost with selling these guarantees, and we are interested in the rate of return on this investment. We define the annualized return on capital (ARC) as

$$\text{ARC} = \frac{1}{t^*} \times \left(\frac{(\text{capital value} + \text{hedge value} - \text{payoff value})\big|_{t^*}}{\text{initial capital requirement}} - 1 \right)$$

This can be regarded as the return on the initial capital per year for the writer of the guarantee. The ARC is similar to the risk adjusted return on capital (RAROC) described in (Jameson 2001), but has been converted to an annualized rate of return to facilitate comparisons between contracts of different durations.

It should be noted that the ARC cannot be thought of as a compounded rate of return. We have chosen this specification of the return on the initial investment since the final value of our position at maturity can be negative, and this cannot be quantified as a compounded rate of return on the (positive) initial investment. In order to facilitate comparisons with compounded rates of return we can define an effective continuously compounded rate

$$r_{\text{eff}} = \frac{1}{t^*} \times \ln(1 + ARC \times t^*)$$

where t^* is the average duration of the contract during the simulation and we use the mean ARC in this calculation. This definition of the effective rate r_{eff} incorporates the fact that, upon selling these contracts the insurer is locked into this position for a duration of time which depends upon the investor's actions. We find that t^* is approximately 6.3 years for the contract with no resets and is approximately 21.2 years for the contract with two resets per year when the investor uses heuristic rules described in this section to determine their use of the reset feature and anti-selective lapsing behaviour. When the investors act optimally, the average duration of the contracts become 6.4 years and 17.9 years respectively.

Table 2 approximately here.

The results in Table 2 indicate that, when no hedging strategy is in place and investors act non-optimally, the expected effective return on a 95% CTE capital requirement is is approximately 9.6% for the guarantee that does not offer any resets, and is about 8.5% for the guarantee which offers two resets per annum. Many segregated funds are currently offered with substantially lower proportional fees. Of course with a lower proportional fee charged to cover the cost of providing the guarantee the return on capital will be reduced. It is interesting to note in Table 2 that, when no hedging strategy is implemented, investor non-optimality does not significantly increase the rate of return on the initial capital investment required by the insurer to satisfy the OSFI guidelines. In a later section, we will find that when a dynamic hedging strategy is in place, investor non-optimality can result in a significant increase in the effective rate of return on capital.

Although the expected return indicates that writing these contracts and leaving them unhedged can be profitable, the capital requirements of $8.65 per hundred dollars of underlying mutual fund for the contract with no resets and $13.46 per hundred dollars for the contract with two resets per annum may be prohibitive. As mentioned above, the OSFI capital requirements for these products can be reduced if appropriate hedging strategies have been put in place. Furthermore, in Figure 1 we see that there is a substantial amount of variability in the ARC, particularly for the contract that offers two resets per annum, with many outcomes generating losses. In the following sections we will investigate the statistical performance of various hedging strategies for segregated fund guarantees.

Hedging Risk Exposure for Segregated Fund Guarantees

The hedging strategies that we will investigate in this paper will incorporate the strengths of both actuarial and modern financial theory approaches. An insurer offering a segregated fund guarantee is exposed to several sources of risk.

For example, due to the mortality benefits offered by many of these contracts, the value of the contract will depend upon the demographic profile of the investor who is purchasing the contract; i.e. female, aged 50 years, non-smoker. If a large number of these contracts have been sold to investors from a similar demographic profile, then we can assume that mortality risk is diversifiable. We can consider hedging an aggregate contract from which a fraction of the investors die during each year with a rate specified by a standard mortality table.

Another source of uncertainty, which may be considered to be diversifiable, is deterministic investor lapsing. Here we may be able to treat the fraction of the investors that withdraw their accounts (for non-optimal reasons) each year as a deterministic function. We would like to emphasize that pricing these contracts under the assumption that investors will act non-optimally may be dangerous and may result in mis-pricings by the insurer. Although the majority of individual investors may not have the expertise to utilize these complex features efficiently, we have heard of incidents where financial planners have assisted their customers in doing so as an additional service.

The insurer is also exposed to risk due to the uncertain movements of the underlying mutual fund since the guarantee will only have positive value to the investor if the mutual fund is below the guarantee level at maturity. It is well known that market risk exposure is not readily diversifiable. In this case techniques from modern financial theory described in Hull (1997) or Wilmott (1998) can be applied.

We denote $V(S, K, U, T, t)$ to be the value of the segregated fund guarantee which depends upon the value of the underlying mutual fund S, the current guarantee level K, the number of reset opportunities used this year U, the current maturity date T, and time t. In this work we consider simple dynamic hedging strategies which create delta-neutral ($\Delta = V_S$) positions for the insurer over brief time intervals. As a result, over short time intervals, the value of the hedged portfolio is immune to changes in the value of the mutual fund. If it is not optimal to utilize the reset feature or lapse then the value V satisfies the partial differential equation (PDE)

$$V_t + (r - r_m - r_e)SV_S + \tfrac{1}{2}\sigma^2 S^2 V_{SS} - rV - R(t)r_e S + M(t)\max(K - S, 0) = 0$$

where $R(t)$ denotes the number of investors remaining in the contract (who have not perished or lapsed) at time t and $M(t)$ is the mortality rate at time t. This equation is very similar to the classical Black-Scholes equation from option pricing theory, but contains two additional terms. The final two terms this equation represent the rate of incoming proportional fees collected from investors remaining in the contract at time t and the rate of payments made to deceased investors at time t respectively. Also, the convection coefficient (in front of the V_S term) is slightly different from that in the classical Black-Scholes model. This is because we have assumed that hedging is performed with a perfectly correlated asset that does not have management fees deducted, and hence the drift rate is affected. As an example, we can consider hedging a guarantee on an index tracking mutual fund by trading index participation units. In a later section we will generalize these techniques to allow for hedging with an imperfectly correlated asset.

If we let $U_{\max}$ denote the maximum number of resets permitted per annum by the segregated fund contract, then if $U < U_{\max}$ there are reset opportunities remaining and V must satisfy the constraint

$$V(S, K, U, T, t) \geq V(S, S, U+1, t+T_{ext}, t)$$

where T_{ext} is the amount that maturity is extended by upon resetting the guarantee level. Effectively, this models the fact that the investors can receive a new guarantee with guarantee level $K = S$ and maturity $T = t + T_{\text{ext}}$, and that one more reset opportunity has been used.

It will be optimal for investors to lapse if the value of the guarantee, net of the proportional fees required to maintain the contract, is more negative than any deferred sales charges that must be paid upon terminating the contract. As a result, the writer cannot allow the value of the hedging position to become negative in anticipation of future incoming fees. We can model optimal investor lapsing by imposing the additional constraint

$$V(S, K, U, T, t) \geq 0.$$

Finally, at maturity, the value of the contract is

$$V(S, K, U, T, T) = R(T) \max(K - S, 0)$$

which states that only investors remaining in the contract at maturity receive the final payoff.

Typically the investor is not charged an initial premium to enter into these contracts. As a result, the fair value for these contracts is determined by the expense ratio, r_e, that makes the value of the contract initially zero. In (Windcliff et al. 2002) we determine the fair proportional fee rate for various contracts and various models for the underlying security. In this work we will assume that the expense ratios are fixed at the levels given in Table 1 and our focus will be to investigate our ability to hedge the risk exposure due to price movements by the underlying security. Other sources of risk, such as interest rate risk and implied volatility risk, basis risk, liquidity risk, etc., will affect the value and hedging of these contracts. The new actuarial reserving guidelines and OSFI's new capital standards require insurers to explicitly provide for these risks if they intend to take credit for hedging strategies. Interested readers are referred to (Windcliff et al. 2001) for a complete and detailed description of the mathematical model and computational techniques used to obtain the hedging strategies for the numerical experiments in this paper.

Statistical Results for Dynamically Hedged Positions

The insurer may wish to implement a hedging strategy for several reasons. First, by implementing a hedging strategy, the downside risk associated with writing these contracts may be reduced. Second, if the insurer implements a hedging strategy, the OSFI capital requirement can be reduced by up to a maximum of 50% of the reduction in the 95% CTE indicated by the proposed hedging strategy. For the numerical results provided in this paper the required capital for a dynamically hedged position will be given by

$$\text{OSFI Capital Requirement} = \text{CTE}_{\text{hedge}} + \tfrac{1}{2}(\text{CTE}_{\text{no hedge}} - \text{CTE}_{\text{hedge}})$$

where the conditional tail expectations are taken with a 95% confidence level.

In Table 3 we provide numerical results when a delta-neutral hedging strategy which is re-balanced 50 times per year (i.e. approximately on a weekly basis) is implemented. Comparing these results with those for unhedged positions contained in Table 2, we see that the 95% CTE is reduced dramatically due to the hedging strategy, resulting in much smaller capital requirements. Of course this reduction in downside comes also comes with lower expected profits from the contract. However, since less capital is required when a hedging strategy is implemented, the rate of return on the initial capital outlay is only moderately affected. For the contract with no resets the rate of return on the initial reserve is approximately 7.6%, whereas for the contract with two resets per annum the return is between 6.3% and 6.9%, depending upon the degree of optimality displayed by the investor in their use of the reset feature.

Table 3 approximately here.

It should be noted that when this dynamic hedging strategy is implemented, non-optimal investor behaviour could now lead to additional profits by the insurer, which was not the case when no hedging strategy was implemented. This indicates that the insurer is not penalized for hedging the worst case situation which assumes that the investor acts optimally as profits will be accrued as non-optimal decisions occur. Although it may presently be safe to assume that the majority of investors will act sub-optimally, it would be dangerous to build this assumption into the pricing and hedging of these products.

Due to the fact that the capital requirement is only reduced by a maximum of 50% of that indicated by the proposed hedging strategy, the effective rate of return on the initial capital outlay by the insurer is quite low. The guidance note issued by OSFI that specifies this maximum capital offset due to hedging (OSFI 2001) states that "as the industry and OSFI gain confidence in implementing such strategies, this limitation will be reviewed." In Table 3 we also provide numerical results in the cases when the capital requirement can be reduced by up to a maximum of 75% and 100% of the reduction in the 95% CTE indicated by the proposed hedging strategy. As expected, in this case the return on capital can improve quite dramatically; particularly when investors act sub-optimally.

In Figure 2(a) we can compare the relative strengths and weaknesses of the hedged position. This figure depicts the ARC for a standard OSFI capital requirement, which allows the initial capital outlay to be reduced by up to a maximum reduction of 50% as a result of the proposed hedging strategy. For the hedged position the initial capital is $7.22 per hundred dollars of underlying mutual fund, while for the unhedged position the initial capital is $13.40 per hundred dollars of underlying mutual fund. In this case we see that for the hedged position, the number of outcomes which generate a loss on the initial capital is negligible; which is not the case for the unhedged position. On the other hand, this reduction in downside risk has come at the expense of upside potential and we see that the hedged position also has relatively fewer outcomes that generate large profits.

Figure 2 approximately here.

Figure 2(b) compares the effects of optimal and heuristic investor behaviour for a hedged position. We see that the profit for the writer increases when the investor does not act optimally. Notice that non-optimal investor behaviour introduces a positive skew in the distribution and the downside risk is not dramatically affected by the heuristic investor behaviour.

Hedging with a Correlated Asset

Classical hedging strategies mitigate the insurer's downside risk by taking a short position in the underlying mutual fund. For obvious reasons, due to regulatory policy this is not possible. If the mutual fund is tracking an actively traded index, then one can use index participation units to accurately hedge risk exposure. The numerical results provided thus far in this paper have assumed this situation. However, it is often the case that the mutual fund is not constructed to closely track an index and hedging must be performed using a basket of securities that closely replicate the performance of the mutual fund. In general, the price movements of this basket of hedging securities will not be perfectly correlated with the underlying mutual fund. Other possible motivations for studying hedging with an imperfectly correlated asset include illiquidity in the underlying asset (Sircar and Papanicolaou 1998) and transaction costs (Wilmott 1998). In this situation, the positions taken by the hedging strategy can affect the price of the underlying asset. Sufficient illiquidity may make hedging with a correlated liquid asset the preferred choice.

In this section we extend the Black-Scholes framework to allow for the pricing and hedging of option contracts when it is not possible to establish a hedging strategy which trades directly in the underlying asset. With the exception of a brief note (Seppi 1999) on minimum variance cross hedging strategies, very little work has been done in the mathematical modelling of hedging strategies in this setting. Given an asset which is correlated with the underlying, we provide a partial differential equation (PDE) which determines a cross-hedging strategy which can be used by the insurer to minimize the variance of the risk exposure to price movements by the underlying asset. The resulting option pricing model includes the Black-Scholes and present expected value models as special cases. Another method that can be applied when hedging with an imperfectly correlated asset formulates the option pricing problem in an incomplete market setting and uses a utility maximization approach (Davis 2000).

Mathematical Model. For expositional simplicity we will develop the model for a hedging with an imperfectly correlated asset in the context of a simple vanilla put option. In particular we ignore exotic features such as mortality benefits, the deferred payment of these contracts through proportional fees, the reset feature and lapsing. The numerical results provided in this paper will be based on a generalized model that incorporates these effects.

Consider an option written on the underlying asset with price given by S_u, which satisfies the stochastic differential equation (SDE)

$$dS_u = \mu_u S_u dt + \sigma_u S_u dz_u .$$

Here μ_u is the drift rate and σ_u is the volatility of this asset and dz_u is an increment from a Wiener process.

If it is not possible to trade directly in the underlying asset S_u, we can try to establish a hedge for this option by trading in another asset with price process S_h, which satisfies

$$dS_h = \mu_h S_h dt + \sigma_h S_h dz_h ,$$

where μ_h is the drift rate and σ_h is the volatility of the asset S_h. The Wiener increment dz_h is correlated with the increment for S_u with $\text{corr}(dz_u , dz_h) = \rho$.

Following standard techniques as described in (Wilmott 1998) we establish a portfolio that contains the option, whose value is given by $V(S_u ,t)$, and a short position of Δ_h shares of the second asset S_h,

$$\Pi = V(S_u , t) - \Delta_h S_h .$$

Using Ito's Lemma we can estimate the change in value of this portfolio over small increments of time. The choice of Δ_h which minimizes the variance of the returns on this portfolio is given by

$$\Delta_h = \rho \frac{S_u}{S_h} \frac{\sigma_u}{\sigma_h} V_{S_u} .$$

We can loosely think of this model as hedging as much of the risk exposure to underlying price movements in light of the basis risk introduced by hedging with a non-perfectly correlated asset.

If this partially hedged portfolio earns the rate of return $\bar{r}$ (which we discuss below) then we find that the value of the option satisfies the PDE

$$V_t + \left(\mu_u - \rho \frac{\sigma_u}{\sigma_h} \left(\mu_h - \bar{r} \right) \right) S_u V_{S_u} + \tfrac{1}{2} \sigma^2 S^2 V_{S_u S_u} - \bar{r} V = 0.$$

In order to determine an appropriate specification for the discounting rate $\bar{r}$ we consider this equation under several special circumstances.

- $\boldsymbol{r} = \pm 1$: In this case we are in a standard Black-Scholes setting and holding Δ_h shares of the hedging asset eliminates all risk from the portfolio Π to leading order. In this case we should discount at the risk-free rate, r.
- $\boldsymbol{r} = 0$: If $\boldsymbol{r} = 0$ then $\Delta_h = 0$ and the portfolio Π contains only a long position in the option. In this case we should discount the portfolio Π by the expected rate of return on the option, r^*, which can be empirically estimated using market option prices and the real-world (P-measure) drift rate.

A simple way to model the discounting rate, $\bar{r} = \bar{r}(\boldsymbol{r})$, which is consistent with the cases described above is to specify

$$\bar{r}(\boldsymbol{r}) = (1 - |\boldsymbol{r}|) r^* + |\boldsymbol{r}| r,$$

where r is the risk-free rate and r^* is the expected rate of return on the option. We remark that when $\boldsymbol{r} = 1$ we recover the Black-Scholes model and when $\boldsymbol{r} = 0$ we recover present expected valuation methods using the real drift rate of the underlying security and a risk-adjusted discounting factor.

Numerical Experiments. In Table 4 we provide estimates for the risk adjusted discounting rate r^* for these contracts. In this work, we cannot estimate r^* using market prices since these exotic long-term options are not traded on exchanges. Instead, we estimate r^* by determining the discounting rate so that the present expected value of these contracts (using the real drift rate for the underlying security) is initially zero. In other words, we determine the discounting rate that the customer must implicitly be using to warrant entering into these contracts.

Table 4 approximately here.

Since the drift rate of the underlying asset is greater than the risk-free rate, taking into account the fees required to maintain the guarantee, the holder of a long position in the guarantee will encounter a loss on average. The results shown in Table 4 are consistent with the findings in (Coval and Shumway 2001) where, using market prices for exchange traded options, the authors find that put options have returns that are both statistically and economically negative. In our case, this refers to the rate of return on the proportional fees paid by the customer to maintain the segregated fund guarantee.

As expected the deferred sales charge has a dramatic impact on the expected rate of return for a contract with no reset features but has very little impact for the contract with two resets per annum. If there are no reset opportunities, the investor will anti-selectively lapse out of the contract if the guarantee becomes out-of-the-money; thereby avoiding the remaining fees required to maintain the guarantee. On the other hand, if the contract offers the customer the ability to reset the guarantee level then anti-selective lapsing does not play as large a role.

In Table 5 we provide results for a partially hedged position when the correlation between the underlying mutual fund and the hedging assets varies between $r = 0$ to $r = 1$. When $r = 0$ we are unable to hedge and the outcome is identical to the unhedged position described in Table 2. When $r = 1$ there is no basis risk and the results are identical to the hedged positions described in Table 3. We see that when the assets are partially correlated the performance of the hedging strategy degrades and the reduction in the required reserve capital is minimal when $r = .9$ and is in fact worse than the unhedged position when $r = .75$. However, the rate of return on the initial capital outlay is not any worse than the case when the hedging asset is perfectly correlated. This is because the portion of the contract that remains unhedged still earns a larger expected profit. Unfortunately, it is probably he case that these strategies will not be useful in practice since we assume that the main advantage of hedging is to reduce the regulatory capital requirements. We do point out that this hedging strategy is in some sense an optimal hedging strategy when basis risk exists in that the position in the hedging asset, Δ_h, was chosen to minimize the variability. As a result, we contend that the management of basis risk is of extreme importance when hedging with a partially correlated asset and should be approached with care.

Table 5 approximately here.

Further, we notice that the re-balancing interval does not significantly affect the performance of the dynamic hedging strategy when using a non-perfectly correlated asset. In Table 5 we see that with a correlation of $r = .9$, adjusting the hedging position 50 times per year only marginally reduces the 95% CTE when compared with hedging 10 times per year, from \$9.60 to \$9.27 per hundred dollars of underlying mutual fund. This is because the majority of the risk is due to the basis risk between the underlying and hedging instruments. In the absence of basis risk it is possible to dramatically improve the performance of the classical Black-Scholes delta-neutral hedging strategy by matching the option's curvature using a gamma hedge. Gamma-neutral hedging strategies reduce the risk exposure to large asset price movements by trading in other option contracts written on the same underlying asset. As noted above, the majority of the risk is due to basis risk and consequently we expect that gamma-neutral strategies would do little to improve the performance when hedging with an imperfectly correlated asset.

Figure 3 plots the profit and loss distributions and distributions of return on capital when hedging with assets which have various degrees of correlation with the underlying mutual fund. We see that even for quite a high correlation of $r = .9$ the distributions are very broad and much of the downside risk is not effectively reduced. In particular, we note that when using a hedging asset with $\rho = .75$, the lower tail of the profit and loss distribution shown in Figure 3(a) is actually thicker than the unhedged case (corresponding to $\rho = 0$). We should point out that many segregated fund guarantees are offered on mutual funds that are actively managed. In this case it may be very difficult to determine a basket of securities that has a very high degree of correlation with the underlying mutual fund.

Figure 3 approximately here.

Conclusion

The decision of whether or not to actively hedge these contracts is a management issue. In essence, hedging can be thought of as constructing insurance in the market for the writer. If one ignores the cost of capital, then hedging reduces the profitability associated with providing these contracts to the customer. However, due to regulatory requirements there is an associated cost with providing these guarantees. In this paper we describe methods for quantifying the return on this initial capital outlay which can be considered when making decision as to whether an institution wishes to continue sale of these contracts and/or whether or not hedging is appropriate. The annualized return on capital (ARC) described in this paper is a form of a risk-adjusted return on capital.

In some cases, due to available resources, management may be forced to establish hedging positions in order to reduce this capital requirement. The results given in this paper indicate that some of the risks involved with offering these contracts can be reduced dramatically using simple dynamic hedging strategies. In fact, as a result of the reduced capital requirements when a hedging strategy is implemented, it may be possible to actually increase the return on the initial capital investment. Currently however, OSFI has taken a conservative position and allows only a reduction of 50% of the indicated reduction in the 95% CTE due to the proposed hedging plan. In this case, the return on capital decreases when hedging is implemented.

In some cases it is not possible to trade directly in the underlying asset. For this situation, we describe a PDE which determines a cross-hedging strategy which trades in a correlated asset and can be used to minimize the variance of the risk exposure of the insurer. The resulting option-pricing model includes the Black-Scholes and present expected valuation models as special cases and the necessary parameters can be estimated using market data. As the correlation between the underlying asset and hedging instrument increases, the variability of the relative hedging error using this hedging model decreases. When hedging with an asset which is not perfectly correlated with the underlying, the majority of the residual risk is due to basis risk between the hedging and underlying instruments. As a result, very frequent re-balancing and more complex gamma-neutral strategies are not very helpful in reducing the variability of the partially hedged position.

To conclude, we point out that the numerical experiments provided in this paper represent a rather conservative contract. In practice many guarantees are offered on more volatile underlying mutual funds and frequently the proportional fees charged to maintain the guarantee are much lower than those used here. In these cases the profitability associated with offering these products can be much lower or can even result in an expected loss. Interested readers are referred to (Windcliff 2002) for additional insights as to why these contracts can be such a liability and how variations in the product design and market parameters such as the volatility and risk-free interest rate affect the cost of providing these guarantees.

Acknowledgment

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

References

Coval, Joshua D. and Tyler. Shumway (2001). Expected option returns. *Journal of Finance 106*(3), 983-1009.

Davis, Mark (2000). Optimal Hedging with Basis Risk. Working paper, Financial and Actuarial Mathematics Group, Technische Universitat Wien, Vienna, Austria.

Hull, John (2000). *Options, Futures and other Derivative Securities, Fourth Edition.* Prentice Hall, New Jersey, United States.

Jameson, Rob (2001, February). Between a RAROC and a hard place. *ERisk*, 1-5.

OSFI (2001, August). *Capital Offset for Segregated Fund Hedging Programs (MCCSR).* Office of the Superintendent of Financial Institutions of Canada.

Seppi, Duane J. (1999). Mathematical Finance: Class notes. Graduate School of Industrial Administration, Carnegie Mellon University.

Sircar, K. Ronnie and George Papanicolaou (1998). General Black-Scholes Models Accounting for Increased Market Volatility from Hedging Strategies. *Applied Math Finance 5*(1), 45-82.

Wilmott, Paul (1998). *Derivatives.* John Wiley and Sons, West Sussex, England.

Windcliff, Heath A., Peter A. Forsyth, and Ken R. Vetzal (2001). Segregated funds: Shout options with maturity extensions. *Insurance: Mathematics & Economics 29*(1), 1-21.

Windcliff, Heath A., Martin K. Le Roux, Peter A. Forsyth, and Ken. R. Vetzal (2002). Understanding the Behaviour and Hedging of Segregated Funds Offering the Reset Feature. To appear in the *North American Actuarial Journal.*

Feature	Description
Investor profile	50 year old Canadian female (from www.soa.org)
Deterministic lapse rate	5% per annum
Optimal Lapsing	Yes, investors will lapse out of the contract if the value of the guarantee becomes less than the value of the remaining fees that will be deducted to maintain the guarantee.
Initial investment	$100
Maturity term	10 years, with a maximum expiry on the investor's 80th birthday.
Resets	Contract 1: No resets. Contract 2: Two resets per year permitted until the investors 70th birthday. Upon reset, the guarantee level is set to the value of the underlying mutual fund and the maturity is extended by 10 years from the reset date.
Mortality benefits	Guarantee is paid out immediately upon the death of the investor.
MER	For both contracts, a proportional fee of $r_m = 1\%$ is allocated to the manager of the underlying mutual fund. In addition, to fund the guarantee portion of these contracts, additional fees are deducted at the following rates. Contract 1: (no resets) $r_e = 50$ b.p. is allocated to fund the guarantee for a total MER of 1.5%. Contract 2: (two resets p.a.) $r_e = 90$ b.p. is allocated to fund the guarantee for a total MER of 1.9%.
DSC	A deferred sales charge is levied upon early redemption using a sliding scale from 5% in the first year to 0% after five years. This fee is paid to the management of the underlying mutual fund and none of it is allocated to the guarantee portion of the contract.
Volatility	$s = 17.5\%$
Risk-free interest rate	$r = 6\%$
Drift rate (before fees)	$m = 10\%$

Table 1: Specification of the guarantee contracts and market information used in the numerical experiments provided in this paper.

		Profit and Loss		Return on Initial Capital		
		Mean	95% CTE	Capital	Mean	r_{eff}
Investor	Contract	($)	($)	($)	ARC	
Heuristic	No resets	1.89	8.65	8.65	13.1%	9.6%
	Two resets p.a.	8.66	13.46	13.46	21.2%	8.5%
Optimal	No resets	1.90	8.72	8.72	13.1%	9.5%
	Two resets p.a.	6.80	13.40	13.40	20.2%	8.5%

Table 2: Statistics for the profit and loss distribution and the return on investment for a 95% CTE capital requirement for an unhedged segregated fund guarantee. The contract with no resets charges a proportional fee of $r_e = 50$ b.p. to fund the guarantee while the contract with two reset opportunities per annum charges a proportional fee of $r_e = 90$ b.p. The heuristic rules used by the investor to determine their use of the reset feature and anti-selective lapsing are described in the accompanying text.

		Profit and Loss		Return on Initial Capital		
		Mean	95% CTE	Capital	Mean	r_{eff}
Investor	Contract	($)	($)	($)	ARC	
Heuristic	No resets	.42	1.02	4.83	9.8%	7.6%
				2.93	11.4%	8.6%
				1.02	19.3%	12.6%
	Two resets p.a.	1.30	.46	6.96	15.7%	6.9%
				3.71	18.6%	7.5%
				.46	62.1%	12.5%
Optimal	No resets	.42	1.02	4.87	9.8%	7.6%
				2.95	11.5%	8.6%
				1.02	19.4%	12.6%
	Two resets p.a.	.28	1.04	7.22	11.8%	6.3%
				4.13	12.4%	6.5%
				1.04	16.1%	7.6%

Table 3: Statistics for the profit and loss distribution and the return on investment for a segregated fund guarantee that is hedged 50 times per year. The three capital amounts shown for each scenario represent a maximum 50% reduction, a 75% reduction and a 100% reduction from the unhedged 95% CTE. The contract with no resets charges a proportional fee of $r_e = 50$ b.p. to fund the guarantee while the contract with two reset opportunities per annum charges a proportional fee of $r_e = 90$ b.p.

Contract	r_e	Lapse penalty	r^*
No resets	50 b.p.	DSC	-14.5%
		No DSC	0.0%
Two resets per annum	90 b.p.	DSC	-7.8%
		No DSC	-7.0%

Table 4: Risk adjusted discounting rates, r^*, for the segregated fund guarantees studied in this paper. We obtained r^* by determining the discounting rate that makes the present value of these contracts zero initially using the real world (P-measure) drift rate for the underlying mutual fund. The deferred sales charge (DSC) used to mitigate investor lapsing utilizes a sliding scale from 5% to 0% during the first five years of the contract.

		Profit and Loss		Return on Initial Capital		
		Mean	95% CTE	Capital	Mean	r_{eff}
Rehedge Frequency	r	($)	($)	($)	ARC	
50 p.a.	1.0	1.30	.46	6.96	15.7%	6.9%
	.9	2.86	9.27	11.37	15.6%	6.9%
	.75	8.61	18.99	18.99	15.8%	6.9%
	.0	8.66	13.46	13.46	21.2%	8.5%
50 p.a.	.9	2.86	9.27			
10 p.a.	.9	2.90	9.60			

Table 5: Performance of hedging strategies which use hedging assets with varying levels of correlation, r, with the underlying mutual fund. The guarantee offers two resets per annum and it is assumed that investors use the heuristic rules as described in the accompanying text for their use of the reset feature and anti-selective lapsing.

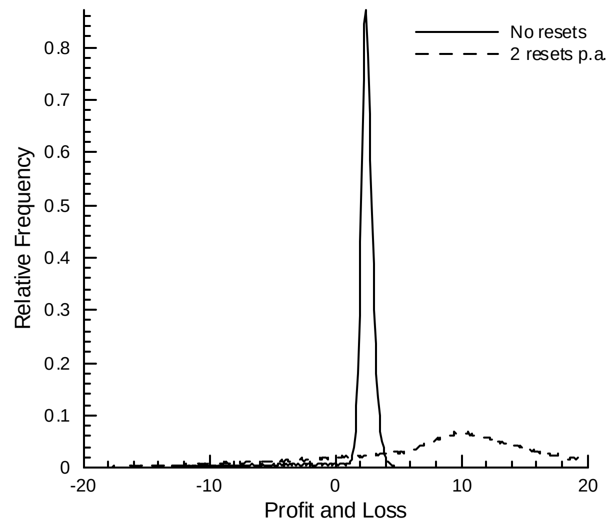

Figure 1(a): The profit and loss distribution for unhedged segregated fund guarantees that offer no resets and two resets per annum. We assume that investors utilize the heuristic rules described in the accompanying text to determine their use of the reset feature and anti-selective lapsing.

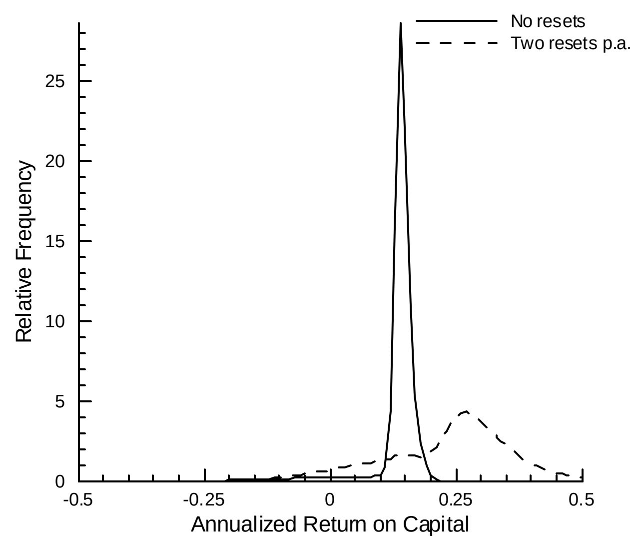

Figure 2(b): The return on investment for a 95% CTE capital requirement for unhedged segregated fund guarantees which offer no resets and two resets per annum. We assume that investors utilize the heuristic rules described in the accompanying text to determine their use of the reset feature and anti-selective lapsing.

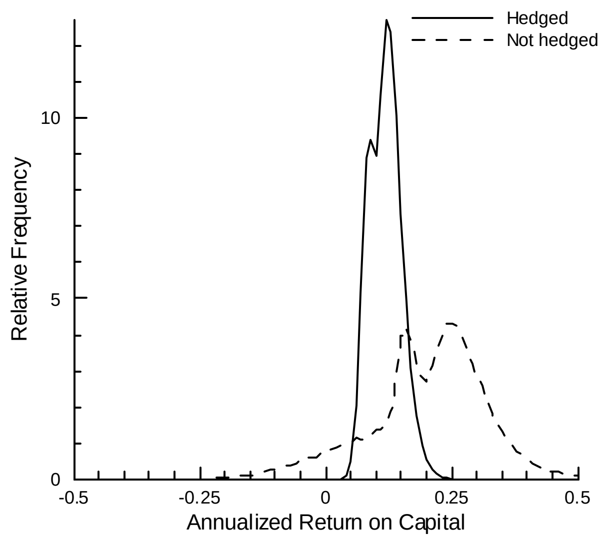

Figure 2(a): Comparison of the return on investment for a hedged position versus an unhedged segregated fund guarantee which offers two resets per annum. The initial capital required for the hedged position is based on a 50% maximum reduction in the 95% CTE from an unhedged position. In this figure, we assume that investors behave optimally.

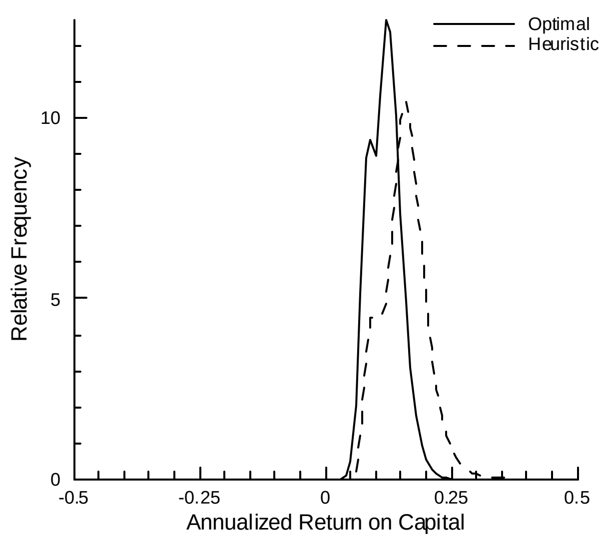

Figure 2(b): Comparison of the return on capital for a hedged segregated fund guarantee which offers two resets per annum when investors act optimally versus heuristic investor behaviour. The initial capital required for the hedged position is based on a 50% maximum reduction in the 95% CTE from an unhedged position.

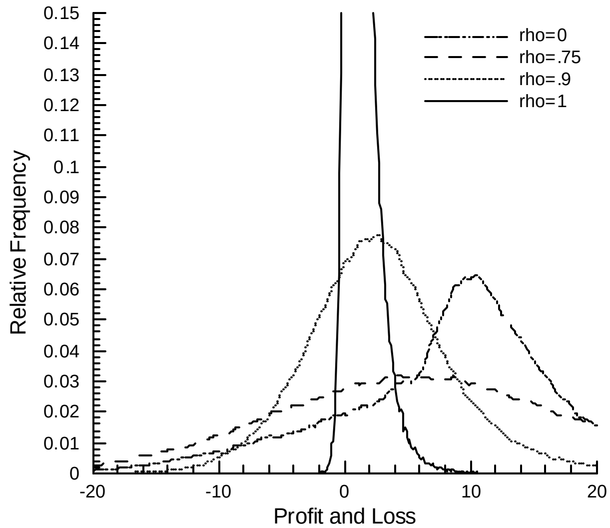

Figure 3(a): The profit and loss distribution for a segregated fund guarantee which offers two resets per annum when hedging using an asset which has correlation r with the underlying mutual fund. We assume that investors utilize the heuristic rules described in the accompanying text to determine their use of the reset feature and anti-selective lapsing.

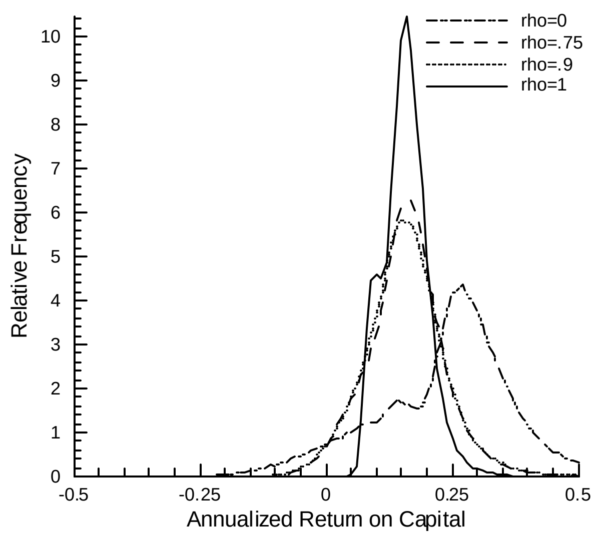

Figure 3(b): The return on capital for a segregated fund guarantee which offers two resets per annum when hedging using an asset which has correlation r with the underlying mutual fund. We assume that investors utilize the heuristic rules described in the accompanying text to determine their use of the reset feature and anti-selective lapsing.