# Hedging Segregated Fund Guarantees

Peter A. Forsyth, Kenneth R. Vetzal and Heath A. Windcliff

**PRC WP 2002-24**

**Pension Research Council Working Paper**

**Pension Research Council**
The Wharton School, University of Pennsylvania
3641 Locust Walk, 304 CPC
Philadelphia, PA 19104-6218
Tel: (215) 898-7620 ● Fax: (215) 898-0310
Email: prc@wharton.upenn.edu
http://prc.wharton.upenn.edu/prc/prc.html

This paper is to be published in *The Pension Challenge: Risk Transfers and Retirement Income Security*. Eds. Olivia S. Mitchell and Kent Smetters. Philadelphia: Pension Research Council Press, forthcoming.

# Hedging Segregated Fund Guarantees

Peter A. Forsyth, Kenneth R. Vetzal and Heath A. Windcliff

Abstract

Segregated funds have become a very popular investment instrument in Canada. Segregated funds are essentially mutual funds which 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. 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 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 which 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.

## Hedging Segregated Fund Guarantees

Segregated funds have become a very popular investment instrument for self-directed pension plans in Canada. Essentially, these are mutual funds that have been augmented with additional insurance features that provide a guarantee on the initial principal invested after a specified time horizon. These investment products offer investors the upside potential of the equity market, while providing a protective floor should the market fall and have been particularly attractive to risk averse investors nearing retirement. The purpose of this chapter is to develop quantitative tools so that institutions can make informed decisions about the risks associated with offering guarantees on mutual funds, and their ability to mitigate this risk through hedging strategies.

Segregated funds 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” that allows investors to increase their guarantee level as the value of the underlying mutual fund goes up. From the viewpoint of the investor, the motivation for the reset feature is easy to see. For example, if the investor starts with an initial investment of $10,000 and the market value of his mutual fund portfolio subsequently increases to $16,000, then the guarantee of recovering his initial principal is unlikely to seem very valuable. This is because, on paper, the investment is currently worth much more than the guarantee level. If the guarantee contract offers a reset provision, the investor can lock in a new guarantee set at the current market value of the account. Typically, the maturity of the new guarantee is ten years from the date at which the guarantee level was set. These contracts also offer features such as mortality benefits, where the guarantee is paid off immediately upon the death of the investor. Recent research has indicated that these guarantees can be very valuable (Windcliff et. al 2002). As a result, the payment for the guarantee is usually amortized over the life of the contract. This introduces additional complexity due to the fact that investors can lapse, and avoid paying any remaining fees, if they deem that the guarantee is not very valuable.

The rationale for financial services firms offering these products is somewhat more subtle. To illustrate, we reconsider our example of an investor who started with an initial investment of $10,000, which has grown to $16,000. If the guarantee does not offer a reset provision, the investor may be better off lapsing out of the guarantee contract in order to avoid paying the remaining proportional fees. Thus, the reset provision can help retain customers. The back end fees used to penalize investors for lapsing often only apply during the first few years of the contract. In fact, in the absence of back end fees, the investor can synthetically create reset opportunities by lapsing and immediately re-entering a new guarantee contract, effectively resetting the guarantee level to the current market value of the account.

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 the guarantees. These capital requirements may be reduced if appropriate hedging strategies have been put in place. This chapter describes hedging strategies that would allow underwriting companies to reduce their risk exposure to these contracts, techniques that incorporate the strengths of both actuarial and financial approaches. We investigate the performance of these types of hedging strategies using stochastic simulation. We also study their impact on the capital requirements used to back these guarantees.

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 short positions in the underlying asset. Since the underlying mutual fund can be under the management of the insurer providing the guarantee, it may not be possible for the underwriting company to take short positions in the underlying mutual fund. One alternative is to hedge using other actively traded securities, such as index participation units, stock index options, or stock index futures contracts, whose behavior is close to that of the underlying mutual fund. This can create basis risk, due to the mismatch between the hedging instrument and the underlying mutual fund. If hedging is performed using short-term derivative contracts such as futures or index options, there will be further risk when the hedging positions are periodically rolled over. Using stochastic simulation we can quantify the risk associated with selling these contracts in a more realistic setting which includes non-optimal investor behavior and basis risk.

We emphasize that the decision of whether or not to hedge these contracts in many situations is a risk-management issue. In some cases, hedging may be necessary to reduce the capital requirements due to regulations. The hedging strategies described here are capable of reducing the downside risk associated with writing such contracts, though they come at an expense, as the expected profit of the hedged position is lower than the expected profit of the unhedged position.

**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 included with many of these contracts is a reset feature. When investors reset, they exchange an existing guarantee for a new 10-year maturity guarantee, set at the current value of the mutual fund. Hence, the reset feature allows them to lock in market gains as the value of the underlying mutual fund increases. The contracts offered in Canada typically allow investors to reset the guarantee level up to a maximum of two or four times per calendar year. This introduces an optimization component to these contracts, where investors must decide when they should reset and lock in at the higher guarantee level.

In addition to the reset option, many other exotic features may be included in segregated fund guarantees. For example, many segregated funds include a death benefit, so that the guarantee is paid out immediately if the investor dies before the maturity date of the guarantee. As the investor ages, these mortality benefits may become more valuable, so resets are often disallowed after the investor's $70^{th}$ birthday. More complex variations of the reset feature can also be introduced, as the investor becomes older. For example, after the investor's $70^{th}$ 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 a proportional fee, $r_m$, allocated to the management of the underlying mutual fund, together with a proportional fee, $r_e$, allocated to finance the guarantee portion of these contracts.

It may be optimal for an 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 a need for liquid assets, investors may sometimes 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 his/her investment during the first several years. Communications with vendors of these contracts suggest that this fee is paid to the underlying mutual fund, and in practice none of it is allocated to finance 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 insurer may not be recovered if the investor lapses prematurely.

The numerical experiments presented in this chapter are based on two contracts depicted in Table 1. The first contract is a simple 10-year maturity guarantee with no reset provisions, while the second contract incorporates a reset feature. These contracts were chosen to be representative of current guarantees offered on mutual funds, and to illustrate the effect of attaching reset provisions to these guarantees. 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; the investor pays for these guarantees by the increased MER as described in Table 1, with the sliding scale DSC to mitigate investor lapsing. Table 1 also describes the key market parameters used in the simulations, such as the assumed risk-free interest rate and volatility of the underlying mutual fund.

*Table 1 here*

Due to the complexity of these contracts, it is difficult to draw general conclusions from individual numerical experiments. The numerical results in this chapter are therefore intended to illustrate the behavior of a realistic contract.[1]

**The Distribution of Returns for Unhedged Positions**

To quantify the risk involved with writing segregated fund guarantees described above, we can investigate the distribution of returns for an unhedged position. At this point, we do not assume optimal investor behavior, but instead we 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 two 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 percent 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 reset opportunities available 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 percent out-of-the-money.

Investigation of the optimal reset region shows that resetting the guarantee when the underlying asset has risen by 15 percent can be a reasonable approximation to the optimal exercise boundary during the first few years of the contract, or during the first few years after a reset has taken place (Windcliff et al. 2002). In fact, this heuristic rule has been adopted by a Canadian Institute of Actuaries task force on segregated funds (CIA 2000) for its assessment of risk management strategies. If investors do not have the ability to reset, 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, 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 (Windcliff et al. 2001).

In the Monte Carlo simulations provided below, it is assumed that the investor makes decisions regarding the reset feature and optimal lapsing 100 times per year, or approximately twice per week.[2] To quantify the risk associated with writing an unhedged segregated fund guarantee, we consider the 95 percent conditional tail expectation (CTE) and the annualized rate of return on an initial capital requirement. The 95 percent CTE is the expected value of the outcomes that lie in the worst case five percent tail. In other words, the 95 percent CTE is the mean value of the worst case outcomes that are ignored by a 95 percent value at risk (VaR) measurement. In comparison with VaR measures, the CTE is much more conservative when setting aside capital for contracts such as segregated fund guarantees, which exhibit a long tail of values that occur with relatively low probability.

Recent regulatory changes from the Canadian regulatory agency OSFI have introduced stricter capital requirements for companies offering these contracts to ensure that sufficient resources are available to back these guarantees. Specifically, if no hedging strategy is put in place, OSFI requires that the insurer set aside the 95 percent CTE in liquid, risk-free instruments. If the insurer implements a hedging strategy, the OSFI capital requirement can be reduced by up to a maximum of 50 percent of the reduction in the 95 percent CTE indicated by the proposed hedging strategy. The capital must be invested in safe, liquid instruments; here we will assume that the capital investment grows at the risk-free rate. It should be pointed out that, for these numerical experiments, 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(a). The P&L for a particular stock price path generated during 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. In this formula, "hedge value" refers to the value of the hedging portfolio at the time the contract is terminated, and "payoff value" refers to the contractual payment made to the investor at the contract termination. 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 percent 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 here*

Table 2 gives statistics for the P&L distributions. It is difficult to draw useful comparisons between these 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. However, one common observation is that the capital requirement, given by the 95 percent CTE, is quite large for any of these contracts. This is the case for both contracts that offer reset provisions, and to a somewhat lesser extent, contracts that offer no resets. We also see that the capital requirement is not substantially reduced by assuming that investors act non-optimally when they make decisions regarding lapsing and when to reset the guarantee level.

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 studying the rate of return on this investment. We define the annualized return on capital (ARC) as:

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

In this formula, "capital value" refers to the value at contract termination of an initial capital outlay of "initial capital" made by the provider of the guarantee. This capital was placed in a risk-free investment to back the guarantee. The quantities "hedge value" and "payoff value" are as described above. This can be regarded as the return on the initial capital per year for the writer of the guarantee.[3]

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 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 define an effective continuously compounded rate, $r_{\text{eff}} = \frac{1}{t^*} \times \ln(1 + \text{ARC} \times t^*)$,

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 $r_{\text{eff}}$ 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 investors use 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 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 percent CTE capital requirement is approximately 9.6 percent for the guarantee that offers no resets, and is about 8.5 percent 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. Note that in Table 2, when no hedging strategy is implemented, investor non-optimality does not significantly increase the rate of return on the insurer's initial capital investment. Below we show 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 requirement may be prohibitive. For the numerical examples provided in this chapter, 8.65 percent of the underlying mutual fund is required for the contract with no resets, and 13.46 percent of the underlying mutual fund value is required for the contract with two resets per annum. 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(b) 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. Next we investigate the statistical performance of various hedging strategies for segregated fund guarantees.

**Hedging Risk Exposure for Segregated Fund Guarantees**

The hedging strategies we investigate 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 (e.g. female, aged 50 years, non-smoker). If a large number of these contracts have been sold to investors from a similar demographic profile, 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 at a rate specified by a standard mortality table.

Another source of uncertainty that 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.[4] 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 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 can be applied (Hull 2000; Wilmott 1998).

Let $V(S,K,U,T,t)$ 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 positions for the insurer over brief time intervals. To create a delta-neutral position, the insurer will purchase

$$\Delta = e^{(r_m + r_e)t} V_S$$

index participation units in order to make the value of the hedged portfolio immune to small changes in the value of the underlying mutual fund over short intervals of time. The exponential factor, which makes this appear different from the typical delta hedge ratio seen in introductory finance texts, is due to the proportional management fees which are deducted from the underlying mutual fund, while it is assumed that no fees are paid on the index participation units.

If it is not optimal to utilize the reset feature or lapse, then the value $V$ satisfies the partial differential equation:

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

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 it contains two additional terms. The final two terms of 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 drift 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. Recall that we are imagining hedging a guarantee on an index tracking mutual fund by trading index participation units. Below, we 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 the guarantee value must satisfy the constraint:

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

where $T_{\text{ext}}$ 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 guarantee payoff.

Typically, the investor is not charged an initial premium or front end load to enter into these contracts. As a result, their fair value is determined by the expense ratio, $r_e$, that makes the value of the contract initially zero. The fair proportional fee rate for various contracts and various models for the underlying fund is described in Windcliff et al. (2002). Here, we assume that the expense ratios are fixed at the levels given in Table 1 and our focus is to investigate the 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.[5]

**Statistical Results for Dynamically Hedged Positions**

An insurer may wish to implement a hedging strategy for several reasons. First, by implementing such a 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 percent of the reduction in the 95 percent CTE indicated by the proposed strategy. For the numerical results provided in this paper the required capital for a dynamically hedged position is given by:

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

where the conditional tail expectations are taken with a 95 percent confidence level. This current capital requirement policy is conservative and does not offer a full credit for insurers that dynamically hedge their positions. We will also study the effect of variations in this capital requirement policy that allow for a full reduction when hedging is implemented.

Table 3 provides numerical results from implementing a delta-neutral hedging strategy which is re-balanced 50 times per year (i.e. approximately on a weekly basis). Comparing these results with those for unhedged positions contained in Table 2, we see that the 95 percent CTE is reduced dramatically due to the hedging strategy, resulting in much smaller capital requirements. Of course, this reduction in downside also comes with lower expected profits from the contract, but 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 percent, whereas for the contract with two resets per annum the return is between 6.3 percent and 6.9 percent, depending upon the degree of optimality displayed by investors in their use of the reset feature.

*Table 3 here*

It should be noted that when this dynamic hedging strategy is implemented, non-optimal investor behaviour could 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, and additional profits accrue as non-optimal decisions occur. Although it may be safe to assume that the majority of investors will act sub-optimally presently, it would be dangerous to build this assumption into the long term pricing and hedging of these products.

The effective rate of return on the initial capital outlay by the insurer is quite low due to the fact that the capital requirement is only reduced by a maximum of 50 percent of that indicated by the proposed hedging strategy. The guidance note issued by OSFI (2001) specifying this maximum capital offset due to hedging 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 percent and 100 percent of the reduction in the 95 percent CTE indicated by the proposed hedging strategy. When it is profitable to offer these contracts, as in the case when investors act sub-optimally, the return on capital can improve quite dramatically when the capital requirements are reduced. As a result, this modified policy may entice more institutions offering these products to implement strategies in order to receive the hedging credit. The risk in allowing for a full capital requirement reduction is that the hedging strategy may prove to be less effective than the model indicates. Below we study the impact of basis risk on the effectiveness of simple dynamic hedging strategies.[6]

Figure 2(a) assesses the relative strengths and weaknesses of the hedged position. This figure depicts the annualized return on capital for a standard OSFI capital requirement, which allows the initial capital outlay to be reduced by up to a maximum reduction of 50 percent 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. We see that for the hedged position the ARC is always positive (to the resolution of the figure). If the ARC is positive, this indicates that the payment made to the investor at the contract's maturity was fully covered by the hedging strategy and capital allocation. In other words, the provider was not required to infuse any additional capital to back the guarantee at the termination of the guarantee contract. On the other hand, the simulation of the unhedged position results in many outcomes where the ARC is negative. It should be noted that hedging, which reduces the downside risk associated with providing these guarantees, also reduces the upside potential. We see that the hedged position also has relatively fewer outcomes that generate large profits when compared with the simulations of the unhedged guarantee.

*Figure 2 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 behavior 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. Since the underlying mutual fund is often under the management of the insurer providing the guarantee, taking short positions directly in the underlying asset is not always possible. We have mentioned above that a guaranteed investment can be viewed as holding the underlying asset and having a put option on that asset. Alternatively, the insurer can back this contract by setting aside a bond with a face value equal to the guarantee level, and dynamically hedging a variation of a 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. On the other hand, now the insurer is exposed to a considerable amount of upside risk if the replication of the call option is not effective.

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. Yet, 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 fund. In general, the price movements of this basket will not be perfectly correlated with the underlying fund. Another possible motivation for studying hedging with a partially correlated asset is illiquidity in the underlying asset (Sircar and Papanicolaou 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. With the exception of a brief note on minimum variance cross hedging strategies (Seppi 1999), very little research has apparently been done in the mathematical modeling of hedging strategies utilizing a partially correlated asset.[7]

In the Appendix, 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. The basic idea is that, given an asset that has price increments which are correlated with the underlying (with correlation coefficient $\rho$), we determine a hedging strategy using this secondary asset. The position held in the secondary hedging asset is chosen to minimize the variance of the partially hedged position. The option pricing model in the Appendix includes, as special cases, the Black-Scholes model, as well as discounted cash flow valuation models which use the actual drift rate and a risk-adjusted discounting rate.

**Numerical Experiments.** Table 4 provides 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 an upper bound on the discounting rate that the customer must implicitly be using to warrant entering into these contracts.

*Table 4 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, on average the holder of a long position in the guarantee will encounter a loss. 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.

Again referring to Table 4, we see that 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 should optimally 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. This indicates that the reset feature may be an effective way for financial institutions to retain customers in these accounts in a rising market.

Table 5 provides results for a partially hedged position when $\rho$, the correlation between the underlying mutual fund and the hedging assets varies between 0 and 1. When $\rho = 0$, the hedging asset and underlying asset are uncorrelated and we are unable to hedge using this asset. In this case, the outcome is identical to that of the unhedged position described in Table 2. When $\rho = 1$, the hedging asset and underlying asset are perfectly correlated and there is no basis risk. In this case, the results are identical to the hedged positions described in Table 3. When the assets are partially correlated the performance of the hedging strategy degrades rapidly. When $\rho = 0.9$, the 95 percent CTE when hedging with the partially correlated asset is \$9.27 compared with \$13.46 for the unhedged position. Consequently, the reduction in the required capital may not be sufficient to warrant attempting to hedge these contracts when basis risk is present. In fact, the 95 percent CTE is even worse when $\rho = .75$, increasing from \$13.46 for the unhedged position to \$18.99 for the position which hedges using the imperfectly correlated asset. This indicates that when hedging with a partially correlated asset, the correlation must be very high. Otherwise, the hedging strategy may, in fact, increase the risk associated with providing these contracts.

*Table 5 here*

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 $\rho = 0.9$, adjusting the hedging position 50 times per year only marginally reduces the 95 percent 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 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 $\rho = 0.9$, the distributions are very broad and much of the downside risk is not effectively reduced. Interestingly, when using a hedging asset with $\rho = 0.75$, the lower tail of the profit and loss distribution shown in Figure 3(a) is actually thicker than the unhedged case (corresponding to $\rho = 0$).

*Figure 3 here*

It must be noted that many segregated fund guarantees are offered on mutual funds that are actively managed. In this case, it may be quite difficult to determine a basket of securities that has a very high degree of correlation with the underlying mutual fund. The hedging strategy described in this section is in some sense an optimal one, in that the position in the hedging asset, $\Delta_h$, is 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.

**Conclusion**

Recent market volatility has spurred interest in attaching guarantees to pension investments. In this chapter, we look at maturity guarantees currently offered on mutual funds in Canada. The guarantees embedded in these products are often much more complex than a simple maturity guarantee that insures the initial investment. In addition to offering mortality benefits, these contracts typically allow the investors to lock in market gains by resetting the guarantee level to the current value of the account. Even if such provisions are not explicitly offered, investors can synthetically create reset opportunities by lapsing and re-entering the contract. Hence, institutions offering even simple maturity guarantees must carefully consider the effects of anti-selective investor lapsing when they are quantifying their risk exposure.

A guarantee is only valuable to the consumer if the financial institution offering the product is able to back the guarantee in the event of a market downturn. In order to ensure solvency, regulatory agencies require that institutions offering these products set aside capital in a risk-free, liquid investment. This introduces an associated cost with providing these guarantees. These capital requirements can be quite onerous, but they can be reduced if a suitable hedging strategy is implemented.

The decision of whether or not to actively hedge guarantees offered on mutual fund investments is a management issue. In essence, hedging can be thought of as constructing insurance in the market for the provider, and our techniques can assist in determining if hedging is appropriate. In particular, we focus on two main questions: how effective is the hedging strategy at removing downside risk, and what is the return on the regulatory capital investment? We found that even a simple delta-neutral hedging strategy was very effective at reducing the downside risk if it was possible to set up a hedging position using the underlying mutual fund, or a perfectly correlated asset (such as hedging an index tracking mutual fund using index participation units). However, in many cases it is not possible to use the underlying mutual fund itself when constructing the hedging portfolio. 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 may not be effective at further reducing the variability of the partially hedged position.

The other main reason why an insurer would consider implementing a hedging strategy is to reduce the regulatory capital requirements for these contracts. Our results indicate that some of the risks involved with offering these contracts, and hence the capital requirements, can be reduced dramatically using simple dynamic hedging strategies. Current regulatory policy in Canada has taken a conservative position and by implementing a hedging strategy, the provider is allowed a maximum 50 percent reduction in the required capital. In this case, the return on the initial capital investment made by the insurer decreases when hedging is implemented. As a result, many institutions offering these products back these guarantees with capital reserves and do not actively hedge their risk exposures to these contracts. There are indications that a full credit for hedging may be granted in the future. In this case, hedging can dramatically reduce the required capital, thereby increasing the return on this initial capital outlay, so more institutions may be inclined to take advantage of this credit.

This work was supported by the Natural Sciences and Engineering Research Council of Canada,the Social Sciences and Humanities Research Council of Canada, and RBC Financial Group. Opinions remain those of the authors.

**Appendix A: Hedging with Basis Risk**

For expositional simplicity, we develop the model for hedging with a partially correlated asset in the context of a simple vanilla put option. In particular, we ignore exotic features associated with segregated fund guarantees such as mortality benefits, the deferred payment of these contracts through proportional fees, the reset feature, and lapsing. Yet, the numerical results provided in this chapter are 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:

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

Here $\mu_u$ is the drift rate and $\sigma_u$ 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 $\mu_h$ is the drift rate and $\sigma_h$ is the volatility of the asset $S_h$. The Wiener increment $dz_h$ is correlated with the increment for $S_u$ with $\mathrm{corr}(dz_u, dz_h) = \rho$.

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 $\Delta_h$ 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 $\Delta_h$ 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 may 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 partial differential equation:

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

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

- $\rho = \pm 1$: In this case we are in a standard Black-Scholes setting and holding $\Delta_h$ shares of the hedging asset eliminates all risk from the portfolio $\Pi$ to leading order. In this case we should discount at the risk-free rate, $r$.
- $\rho = 0$: If $\rho = 0$ then $\Delta_h = 0$ and the portfolio $\Pi$ consists only of a long position in the option. In this case we should discount the portfolio $\Pi$ 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, $\bar{r} = \bar{r}(\rho)$, which is consistent with the cases described above is to specify

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

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

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

| Feature | Description |
|---|---|
| Investor profile | 50 year old Canadian female (from www.soa.org) |
| Deterministic lapse rate | 5% per annum |
| Optimal Lapsing | 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.<br>Contract 2: Two resets per year permitted until the investor's 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 finance the guarantee portion of these contracts, additional fees are deducted at the following rates.<br>Contract 1: (no resets) $r_e = 50$ b.p. is allocated to finance the guarantee for a total MER of 1.5%.<br>Contract 2: (two resets p.a.) $r_e = 90$ b.p. is allocated to finance 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 | $\sigma = 17.5\%$ |
| Risk-free interest rate | $r = 6\%$ |
| Drift rate (before fees) | $\mu = 10\%$ |

Source: Authors' computations.

**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**

| | | Profit and Loss | | Return on Initial Capital | | |
|---|---|---|---|---|---|---|
| | | Mean | 95% CTE | Capital | Mean | $r_{\text{eff}}$ |
| 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% |

Source: Authors' computations.
Note: The contract with no resets charges a proportional fee of $r_e = 50$ b.p. to finance 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 investors to determine their use of the reset feature and anti-selective lapsing are described in the accompanying text.

**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**

| | | Profit and Loss | | Return on Initial Capital | | |
|---|---|---|---|---|---|---|
| | | Mean | 95% CTE | Capital | Mean | $r_{\text{eff}}$ |
| 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% |

Source: Authors' computations.
Note: 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 finance the guarantee while the contract with two reset opportunities per annum charges a proportional fee of $r_e = 90$ b.p.

**Table 4. Risk adjusted discounting rates, $r^*$, for the segregated fund guarantees studied in this paper**

| 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% |

Source: Authors' computations.
Note: The discounting rate, $r^*$, was obtained 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.

**Table 5: Performance of hedging strategies which use hedging assets with varying levels of correlation, $\rho$, with the underlying mutual fund**

| | | Profit and Loss | | Return on Initial Capital | | |
|---|---|---|---|---|---|---|
| Rehedge Frequency | $\rho$ | Mean ($) | 95% CTE ($) | Capital ($) | Mean ARC | $r_{\text{eff}}$ |
| 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 | | | |

Source: Authors' computations.
Note: 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.**

Source: Authors' computations.

Note: 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.

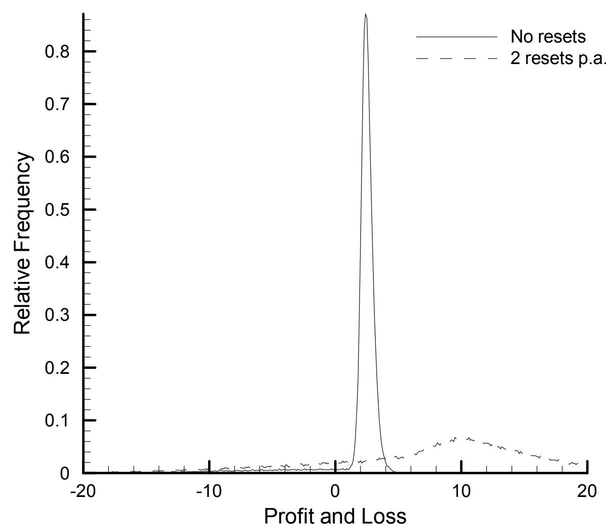

**Figure 1(b): The return on investment for a 95% CTE capital requirement for unhedged segregated fund guarantees which offer no resets and two resets per annum.**

Source: Authors' computations.

Note: Model assumes that investors utilize the heuristic rules described in the accompanying text to determine their use of the reset feature and anti-selective lapsing.

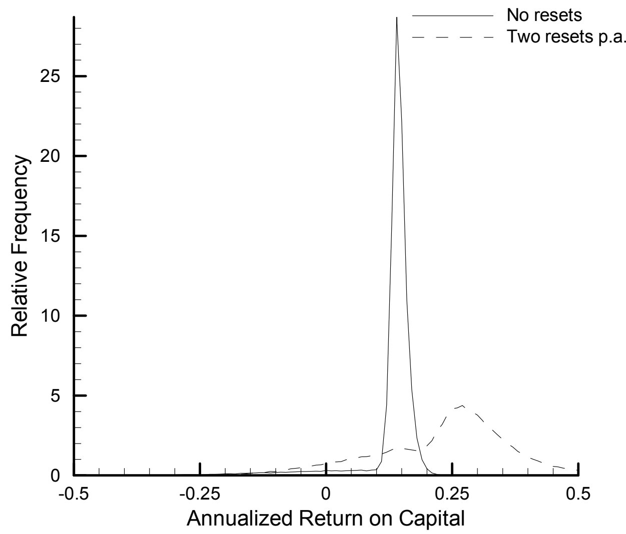

**Figure 2(a): Comparison of the return on investment for a hedged position versus an unhedged segregated fund guarantee which offers two resets per annum.**

Source: Authors' computations**.**

Note: The initial capital required for the hedged position is based on a 50% maximum reduction in the 95% CTE from an unhedged position. Figure assumes that investors behave optimally.

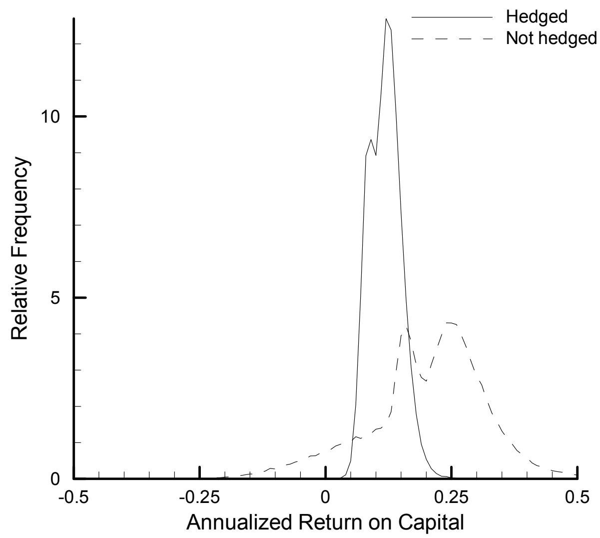

**Figure 2(b): Comparison of the return on capital for a hedged segregated fund guarantee which offers two resets per annum when investors act optimally versus heuristic investor behavior.**

Source: Authors' computations**.**

Note: The initial capital required for the hedged position is based on a 50% maximum reduction in the 95% CTE from an unhedged position.

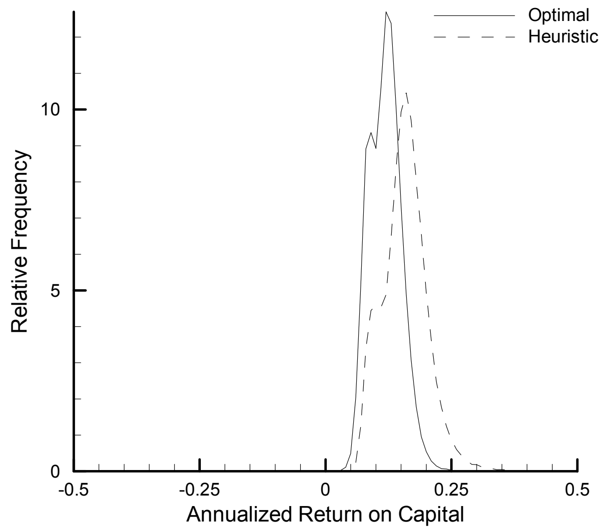

**Figure 3(a): The profit and loss distribution for a segregated fund guarantee which offers two resets per annum when hedging using an asset which has correlation $\rho$ with the underlying mutual fund.**

Source: Authors' computations**.**

Note: Model assumes that investors utilize the heuristic rules described in the accompanying text to determine their use of the reset feature and anti-selective lapsing.

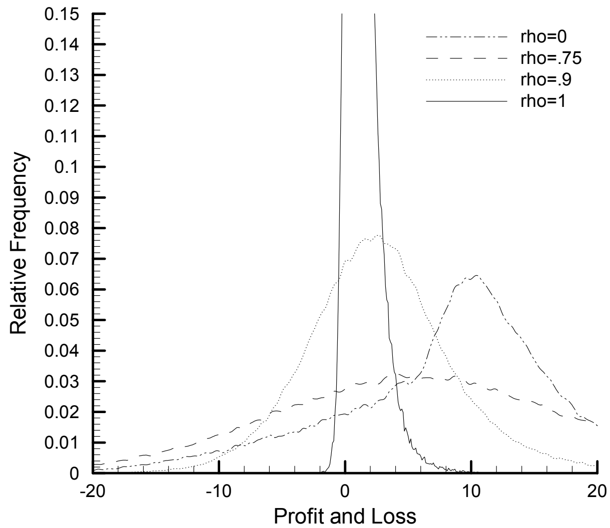

**Figure 3(b): The return on capital for a segregated fund guarantee which offers two resets per annum when hedging using an asset which has correlation $\rho$ with the underlying mutual fund.**

Source: Authors' computations**.**

Note: 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.

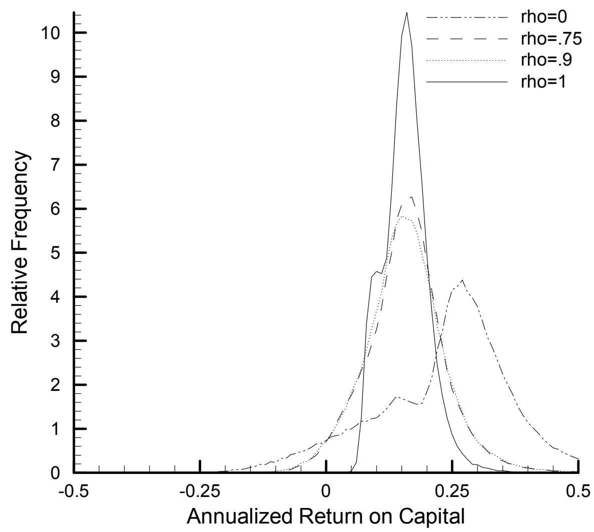

**References**

Canadian Institute of Actuaries. 2000. “Report of the CIA Task Force on Segregated Fund Investment Guarantees.” <www.actuaries.ca/publications/2000/20020e.pdf>.

Coval, Joshua D. and Tyler Shumway. 2001. “Expected Option Returns.” *Journal of Finance 56*(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.* New Jersey: Prentice Hall.

Jameson, Rob. 2001. “Between a RAROC and a Hard Place.” *ERisk,* February: 1-5.

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

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 Kenneth 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 Kenneth R. Vetzal. 2002. “Understanding the Behaviour and Hedging of Segregated Funds Offering the Reset Feature.” *North American Actuarial Journal.*

**Endnotes**

[1] Small changes in the contractual details or variations in the market settings such as the volatility and risk-free interest rate may affect the pricing and hedging of these contracts. Interested readers are referred to Windcliff et al. (2002) for a discussion of the effects of these and other parameters such as the level of investor optimality used in making reset decisions on the valuation of segregated fund guarantees.

[2] Numerical experiments indicate that more frequent exercise decisions by the investor do not appreciably affect results.

[3] The ARC is similar to the risk adjusted return on capital (RAROC) described in Jameson (2001), but it has been converted to an annualized rate of return to facilitate comparisons between contracts of different durations.

[4] We 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.

[5] Interested readers are referred to Windcliff et al. (2001) for a detailed description of the mathematical model and computational techniques used to obtain the hedging strategies for the numerical experiments in this paper.

[6] For the impact of other modeling assumptions such as volatility, interest rates, investor profile, as well as the impact of variations in the product design, interested readers are referred to Windcliff et. al (2002).

[7] Another method that can be applied when hedging with a partially correlated asset formulates the option pricing problem in an incomplete market setting and uses a utility maximization approach (Davis 2000).