Seeing Beyond Risk Canadian Institute of Actuaries Institut canadien des actuaires Voir au-delà du risque

Educational Note

Reflection of Hedging in Segregated Fund Valuation

Committee on Life Insurance Financial Reporting

May 2012

Document 212027

Ce document est disponible en français

Members should be familiar with educational notes. Educational notes describe but do not recommend practice in illustrative situations. They do not constitute Standards of Practice and are, therefore, not binding. They are, however, intended to illustrate the application (but not necessarily the only application) of the Standards of Practice, so there should be no conflict between them. They are intended to assist actuaries in applying Standards of Practice in respect of specific matters. Responsibility for the manner of application of Standards of Practice in specific circumstances remains that of the members in the life insurance practice area.

Memorandum

To: All Fellows, Affiliates, Associates and Correspondents of the Canadian Institute of Actuaries

From: Phil Rivard, Chair
Practice Council

Edward Gibson, Chair
Committee on Life Insurance Financial Reporting

Mike Schofield, Chair
Working Group on the Reflection of Hedging in Segregated Fund Valuation

Date: May 9, 2012

Subject: **Educational Note: Reflection of Hedging in Segregated Fund Valuation**

In accordance with the Canadian Institute of Actuaries' Policy on Due Process for the Approval of Guidance Material Other than Standards of Practice, this Educational Note has been prepared by the Working Group on the Reflection of Hedging in Segregated Fund Valuation, and has received final approval for distribution by the Practice Council on May 8, 2012.

As outlined in subsection 1220 of the Standards of Practice, "*The actuary should be familiar with relevant Educational Notes and other designated educational material*." That subsection explains further that a "practice which the Educational Notes describe for a situation is not necessarily the only accepted practice for that situation and is not necessarily accepted actuarial practice for a different situation." As well, "Educational Notes are intended to illustrate the application (but not necessarily the only application) of the Standards, so there should be no conflict between them."

The working group has solicited input from the Committee on Life Insurance Financial Reporting, the Committee on the Appointed/Valuation Actuary and the Committee on Investment Practice.

The members of the working group are

Robert Berendsen	Frédéric Kibrité
Steve Bocking	Pierre-Laurence Marchand
Lisa Forbes	Peter Phillips
Alexis Gerbeau	Mike Schofield
Brian Fortune	Qi Sun
Lynn Guo	Dean Stamp

Should you have any queries or comments regarding this Educational Note, please contact Mike Schofield at Mike.Schofield@sunlife.com.

PR, EG, MS

TABLE OF CONTENTS

1. PURPOSE/SUMMARY

The August 2010 Report of the Task Force on Segregated Fund Liability and Capital Methodologies (document 210053) recommended the creation of one or more working groups to develop additional guidance for how to perform a CALM valuation of segregated fund investment guarantees when such guarantees are fully or partially hedged. This Working Group on the Reflection of Hedging in Segregated Fund Valuation was assigned the mandate of

> providing guidance for the use of approximation methods to account for hedging in the calculation of insurance contract liabilities, and
>
> providing guidance with respect to potential hedging weaknesses that would be reflected in insurance contract liabilities.

Throughout this Educational Note reference is made to the following CIA documents,

> Report: CIA Task Force on Segregated Fund Investment Guarantees, March 2002 (document 202012), hereinafter "the 2002 task force report",
>
> Educational Note: Considerations in the Valuation of Segregated Fund Products, November 2007 (document 207109), hereinafter "the 2007 Educational Note", and
>
> Report: Report of the Task Force on Segregated Fund Liability and Capital Methodologies, August 2010 (document 210053), hereinafter "the 2010 task force report".

The approximation methods investigated by this working group are meant to be approximations to the CALM framework. Specifically, we investigated approximation methods that approximate the first-principles-based stochastic-on-stochastic methodology.

This first portion of this Educational Note focuses on hedging and the inherent risks to be considered in the valuation. Later sections focus on methods for reflecting hedging within the CALM framework:

> first-principles stochastic-on-stochastic method,
>
> adapted risk-neutral method,
>
> stochastic-on-stochastic with hedge asset proxy method,
>
> hedge cost method, and
>
> proxy function methods.

The appendix gives some explanatory examples to aid in the understanding of the various methods. Section 5.7 contains a summary of the pros and cons of the different methods.

This Educational Note does not formally endorse any one approximation method, but it does express concerns with some approximation methods. The actuary is reminded of Standards of Practice subsection 1510, which provides general guidance on the use of approximations, and the 2006 Educational Note on Approximations to Canadian Asset Liability Method.

2. DEFINITIONS

Delta – The first-order sensitivity of the investment benchmark to the change in equity markets or currency markets. Delta is typically broken down into partial deltas with respect to each equity or currency market.

Dynamic Hedging – Hedging using a portfolio of hedge instruments that are frequently rebalanced. The instruments within a dynamic hedge portfolio typically include short-dated instruments, including derivative instruments, which do not match the tenor of the investment benchmark and require rebalancing as market prices move and as time passes. Longer-dated instruments can be used as part of a dynamic hedging strategy, but would not typically be the instruments traded in the regular rebalancing.

Gamma – The second-order sensitivity of the investment benchmark to the change in equity markets or currency markets.

Greeks – The first or higher order sensitivities of the investment benchmark to movements in various market parameters (e.g., delta, gamma, rho, etc).

Hedge Effectiveness – There are two types of hedge effectiveness discussed in this Educational Note. The narrower definition is the effectiveness versus the investment benchmark and how much slippage occurs versus what you are attempting to hedge. The wider definition is the income statement hedge effectiveness and refers to the ability of the hedge program to mitigate overall earnings volatility.

Hedge Policy – A company's hedge policy articulates the investment strategy, objectives, goals, limits, responsibilities, measurement and monitoring of the hedging program. This would include defining what is hedged (i.e., target hedge = what Greeks of the investment benchmark are hedged) and how the effectiveness is measured and monitored.

Historical Volatility – The realized volatility of a market index or financial instrument, typically derived using standard deviation of returns.

Implied Volatility – The volatility used in a pricing model that reproduces the current market price of an option.

Investment Benchmark – The investment benchmark is the representation of the liability that the investment/hedging professionals use to trade or balance against. It is defined by the selection of a valuation/measurement framework (e.g., real-world vs. risk-neutral vs. some accounting measurement) and the scope of cash flows included in the valuation (e.g., whole segregated fund contract vs. guarantee costs and guarantee fees only vs. guarantee costs only vs. specific guarantee). The investment benchmark is sensitivity tested in order to calculate the Greeks. The hedge policy will articulate the Greeks actually being hedged (i.e., the target hedge).

Node – A given time step on a given real-world scenario path. At each node, we (1) calculate the hedge gain or loss from the prior time step, (2) recalculate Greeks for the investment benchmark, and (3) rebalance the hedge portfolio using the Greeks. Approximations may be used for portions or all of these calculations.

Real-world Framework – A framework using a scenario set that is representative of the future returns and variability of returns. Information can be derived from the distribution of returns or values and hence allows for CTE and percentile-type measurement.

Rho – The first-order sensitivity of the investment benchmark to the change in interest rates. Rho is often broken down into partial or key-rate rhos, which are sensitivities to specific portions of the yield curve.

Risk-neutral Framework – A framework that uses market variables (risk-free rates, swaps and equity option prices) to calibrate stochastic scenarios to reproduce market-observed values.

Adaptations of this framework are used in this Educational Note. The terms “risk-neutral” and “market consistent” are used interchangeably in this Educational Note.

Static Hedging – Hedging using a portfolio of hedge instruments that are not rebalanced over time. Static hedging is sometimes used to mitigate the risk exposure of a closed block of business, or simply to provide a bulk offset to market risk in the segregated fund portfolio. The hedges within a static portfolio typically have long tenors to match the long tenor of the liability portfolio. This length tends to minimize the drift that occurs between the liability sensitivities and the hedge sensitivities as time passes.

Target Hedge – A hypothetical portfolio of hedge instruments whose value changes exactly the same way as does the investment benchmark with respect to the hedged Greeks, and does not change value with respect to non-hedged Greeks. For example, if a company hedges delta and rho, then the target hedge’s value should mirror changes in the investment benchmark due to equity market moves and changes in interest rates, but should not change value due to a change in market implied volatility levels, even if such change does affect the value of the investment benchmark.

Vega – The first-order sensitivity of the investment benchmark to the change in the level of volatility.

3. APPROPRIATENESS OF APPROXIMATION METHODS TO CALM

Nested simulations make the first-principles stochastic-on-stochastic (SOS) method (described in section 6.1) very time-consuming and computationally demanding. It is, therefore, likely that many actuaries will prefer to use an approximation method to determine the insurance contract liabilities under CALM. The actuary is reminded that other than the first-principles stochastic-on-stochastic method, the methods described in this Educational Note are approximations to the first-principles CALM valuation.

4. HEDGING IN THE CONTEXT OF CALM VALUATIONS

4.1 Static vs. Dynamic Hedging

Hedging is a form of risk mitigation. In the context of this Educational Note, hedging refers to an insurer’s actions to mitigate its exposure to the financial market risks embedded in the investment guarantees of its segregated fund portfolio. Hedging is generally achieved by entering into financial transactions (often involving derivative instruments) that have the opposite sensitivity to changes in market factors than the segregated fund guarantees, such that when financial markets move, the effect on the value of the segregated fund guarantees is (largely) offset by the change in value on the hedge instruments.

We can classify hedges into two broad categories, static hedges and dynamic hedges. A static hedge refers to a hedge portfolio that does not involve rebalancing. Static hedges are sometimes used to mitigate the risk exposures of a closed block of business, and would typically employ long-dated instruments that attempt to mimic the current and expected future sensitivities of the liabilities to certain market factors. There is no attempt to match the exposures tightly or to rebalance the hedge over time. Dynamic hedges, on the other hand, refer to a hedge that is regularly or dynamically rebalanced. Dynamic hedges can employ both short-dated and long-dated instruments and aim more closely to match/offset the current/short-term sensitivities of the liabilities to certain market factors. The dynamic hedge is frequently rebalanced (for example, on a daily or weekly basis) to re-establish the match between the evolving market sensitivities of the liabilities and the portfolio of hedge instruments.

A company's hedge policy is influenced by its tolerance for risk. A hedge policy typically addresses the risks the company faces, how the risks are to be measured, which of the risks are hedged, how much of each risk is targeted to be hedged, hedge target mismatch tolerances, the types of instruments that can be used to build the hedge, etc. The hedge policy may be specific to the segregated fund product line or one may consider hedging more broadly in a total company context. Only hedging specifically supporting the segregated fund product cash flows would be included within the segregated fund CALM valuation.

If the hedging strategy is a static one, involving no future hedge rebalancing, then the modelling of the existing hedge instruments is no more difficult than the modelling of the segregated fund liabilities. Therefore, this Educational Note focuses on the case where dynamic hedges are in place, though some of the considerations discussed are relevant also to the modelling of static hedges (e.g., basis risk).

When hedging market risks using a dynamic hedging strategy, the hedge policy typically involves a description of which Greeks are to be hedged (i.e., target hedge), where we use the term "Greeks" (delta, rho, etc.) loosely to refer to sensitivities of the item (i.e., investment benchmark) being measured (whether on a market-consistent basis or otherwise) to specific market risk factors (equity market, interest rates, etc.). The hedge policy may be implemented in practice by defining an investment benchmark or liability that the investment area will hedge against.

The investment benchmark typically would be articulated in the hedge policy. The benchmark could be a best estimate liability, or liability with margins for adverse deviations included. The benchmark could be based on a risk-neutral liability or the accounting liability. Note that using the CALM liability as the investment benchmark can result in a circular calculation. This is rarely done in practice and, consequently, is not explored further within this Educational Note. The investment benchmark is also defined by the scope of the cash flows that are hedged (e.g., benefits and fees, benefits only, specific type of benefit, etc.). The chosen hedging strategy will affect the appropriateness of the various approximation methods.

For further details on hedging, please refer to section 3 of the 2010 task force report.

4.2 First-principles Application of CALM with a Dynamic Hedging Program

The remainder of this Educational Note discusses approximation methods for reflecting the impact of a dynamic hedging program in a CALM valuation. Before discussing approximation methods, we examine the exact calculation. The following is adapted from the 2010 task force report.

> When a hedging program is in place, an exact application of CALM would consist of the following steps.
>
> 1) Generate stochastic scenarios of market variables such as investment returns and interest rates using a model under the real-world measure.
> 2) For each scenario,
>   a. project liability cash flows over the term of the liabilities using actuarial assumptions that include MfADs,

b. at each time step, calculate the Greeks (at least those being hedged as articulated in the hedge policy) or sensitivities of the investment benchmark (e.g., if the investment benchmark was a risk-neutral liability a series of risk-neutral scenarios would be shocked in order to determine Greeks),
c. using the information from step b., project the rebalancing of the hedge portfolio and the resulting hedge portfolio cash flows, and
d. perform a roll-forward CALM cash flow test to determine the amount of required assets which reduce to zero at the last liability cash flow, taking into account the cash flows from the hedge portfolio calculated in step c.

3) Calculate the CTE (level 60% to 80%) of the value of required assets.

The insurance contract liability for the guarantees is set to the CTE calculated in step 3) adjusted for any unamortized amortized acquisition expense (AAE). This adjustment, and the revenue included in the liability cash flows in step 2) a., depend on whether the whole contract or the bifurcated approach is adopted.

We examine the first-principles stochastic-on-stochastic (SOS) application more fully in section 6. The remainder of section 4 is relevant regardless of the method used for performing the CALM valuation.

4.3 Hedge Effectiveness

The term "hedge effectiveness" is often employed when describing the performance of hedge programs, but the term has been used with varying meanings. In general terms, hedge effectiveness refers to how well the hedge performs at narrowing the range of financial outcomes. A high value (close to 100%) would indicate that it is very successful. However, the benchmark against which the hedge is measured has a significant effect on the metric's value. To illustrate the point, we include two definitions of hedge effectiveness below, but acknowledge that there may well be more definitions in use. We caution readers to ensure that they understand what definition is used at a given company. Regarding the two definitions included here, we note that the hedge effectiveness described in 4.3.1 will invariably yield a higher (closer to 100%) hedge effectiveness metric than the definition provided in 4.3.2.

Regardless of definition, determining and understanding the sources of hedge ineffectiveness can aid in achieving a full understanding of the functioning of a hedge program and in assessing the appropriateness of modelling assumptions. In particular, it is useful to compare the hedge effectiveness being modelled to the hedge effectiveness being experienced in real life, and to ensure that the model does not overstate the hedge effectiveness actually achieved in practice. In the various methods described in this document, some will be useful in helping to quantify aspects of hedge ineffectiveness, while others will not.

4.3.1 Hedge Effectiveness vs. the Target Hedge

This is a narrow but frequently-used definition of hedge effectiveness. It refers to the ability or effectiveness of a hedge program to eliminate the specific risks it aims to mitigate, typically referring to the slippage in the hedge's investment performance against what you are attempting to hedge (i.e., the target hedge). Even with this narrow definition, all hedging programs have some degree of hedge ineffectiveness. The actuary is encouraged to monitor the effectiveness against the target hedge to ensure the valuation does not assume effectiveness that is better than experienced. Monitoring of the effectiveness may entail an attribution analysis of the change in the investment benchmark. The change in the investment benchmark may be decomposed into target hedge changes and other changes (unhedged Greeks, non-economic experience changes, etc.). This is often necessary to measure the investment/hedge department's performance as they would normally be held accountable for changes in the target hedge only.

It is often the case that a hedging program will not hedge all the economic risks associated with the investment benchmark (in this case the investment benchmark differs from the target hedge). For example, some companies do not hedge vega risk, while others do not hedge bond fund risk. In this context the actuary is encouraged to examine carefully how hedge effectiveness is measured in practice, paying particular attention to how one defines the change in the investment benchmark due to hedged and unhedged risks, so as not to overstate or understate hedge effectiveness, nor to introduce bias when apportioning the change due to hedged and unhedged risks.

In addition to having an appropriate best estimate for the effectiveness of the hedging program, the actuary would incorporate a margin for adverse deviation. Methods of reflecting the hedge ineffectiveness are described later in this Educational Note.

4.3.2 Income Statement Hedge Effectiveness

This is a broad definition of hedge effectiveness. It refers to the ability or effectiveness of a hedge program to reduce income statement volatility, i.e., to reduce the magnitude of deviations from expected earnings. Earnings volatility exists when the hedge assets change differently than does the CALM liability. With such a broad definition, it would be clear that no segregated fund hedge program is able to attain perfect income statement hedge effectiveness.

There are many risks, costs or modelling decisions that can create deviations between assets and CALM liabilities, and hence hinder income statement hedge effectiveness. Some of these risks are highlighted in section 5.2 below. Note that the list is not all-encompassing, so there would be other sources that hinder income statement hedge effectiveness. The complexities of the liabilities themselves will contribute to hedge ineffectiveness. Determining and understanding the sources of the income statement hedge ineffectiveness can aid in achieving a full understanding of the functioning of a hedge program and in assessing the appropriateness of modelling assumptions.

4.4 Risks and Costs in a Hedging Program to Reflect in the Valuation

While the existence of a hedge program reduces a company's exposure to market risks, it also exposes the company to additional risks, described below. These additional risks contribute to the hedge ineffectiveness. The actuary would ensure that these risks and costs are accounted for within the valuation. Some of the methods described in this note can be used to quantify some of the risks described in this section, while other methods do not allow for quantification, and techniques as listed in section 5.1 would be required.

4.4.1 Basis Risk

Most segregated fund products offer managed fund investment options, i.e., non-indexed funds, whose objective is to outperform a benchmark, a specific market index or a combination of market indices. In traditional segregated fund valuations, fund mapping is used to determine the optimal combination of market indices that closely replicates the performance of the non-indexed or managed funds. The effect of active fund management, i.e., the difference between the market index returns and the actual fund returns, is lost in this fund mapping exercise. In the absence of a hedge, this fund mapping practice is sufficient since it focuses the valuation on appropriately capturing the full range of fund value returns and hence the full range of guarantee costs.

Hedging, on the other hand, is typically performed using derivative instruments that derive their value from market indices, not from the value of managed funds. In addition, hedge instruments themselves can slip from the indices they are tracking and index funds can slip from their benchmark indices.

Basis risk is this imperfect alignment between returns on managed funds and those on hedge instruments. When hedging, the hedge instruments may be able to significantly offset the effect of changes in market indices on the liabilities, but the hedge generally cannot offset the effect of the actively managed fund’s performance versus the mapped market indices. In addition, the hedge vehicle slippage versus the market indices also contributes to gains or losses.

Fund mapping studies and analysis would be performed in order to help quantify the magnitude of the basis risk inherent in the hedging of the product. Once quantified, this risk may be reflected in the valuation as described in section 5.

4.4.2 Liquidity Risk, Bid-Ask Spread and Market Effect Costs

Hedging programs can create significant liquidity risk whether using put options, rolling futures positions or transacting in swaps. Liquidity risk is caused by a widening of the spreads on derivatives that increases the transaction costs required to re-balance the hedge portfolio. In the extreme case, it can also be realized by the impossibility of trading hedge instruments because the markets have been intentionally closed or because the spreads are so excessive that no trading is practical. This extreme case may be more appropriately covered by capital requirements than by liabilities. In short, the liquidity risk is related to the non-commission-type cost of the transactions required to rebalance the portfolio. The actuary would monitor the liquidity of hedge instruments as part of the hedge program measurement.

4.4.3 Transaction Costs and Commissions

Hedging programs can have significant amounts of trading which will generate transaction costs, which would be included in the CALM valuation. Similar studying of experience as is done with assets supporting insurance liabilities is appropriate.

4.4.4 Counterparty Risk

Hedging programs may involve additional counterparty risk resulting from derivatives above that inherent in other product lines. The risk would be dependent on the type of derivatives used within the hedge program. The actuary would consider the Standards of Practice related to credit risk.

4.4.5 Volatility Risk

If the dynamic hedging strategy utilizes material amounts of options or instruments with embedded options, then future market-implied volatility becomes an important assumption that will affect the cost of hedging.

The interplay between modelled time-zero volatility and simulated forward volatility can have a dramatic effect on the simulated economics and risk profile of a dynamic hedging program or strategy. For example, many hedge programs are exposed to vega risk. If, over time, in practice or in a laboratory simulation, volatility levels fall sharply, such hedge programs should make money (or vice versa). In this context the actuary is encouraged to consider the important relationship between time-zero volatility and simulated forward volatility to assess properly the material effect that volatility modelling choices can make on expected results and risk profiles.

4.4.6 Risks Intentionally Not Hedged

A hedging strategy may intentionally not hedge some risk types (e.g., volatility risk, portions of interest rate or equity risk), certain elements of the liability cash flows such as fee revenue based on account value, specific funds (e.g., those which do not map well to any market indices), specific benefits (e.g., death benefit not hedged), or margins for adverse deviation (i.e., hedge best estimate liability only).

A bifurcation of the liability may be an appropriate option. As an example, in the case of unhedged fee revenue, bifurcating the fees between unhedged fees and hedged fees and calculating the unhedged portions using a first-principles CALM approach that does not reflect hedging may be appropriate. There may be other bifurcations that are possible. Care would be taken in this type of bifurcation to avoid disconnecting insurance contract feature interactions.

4.4.7 Risks Not Explicitly Modelled

There may be risks that are not explicitly captured in the modelling. These can exist because they are intentionally not modelled or because approximations have been used to simplify the modelling or the valuation method itself cannot capture the risks. These risks would still need to be included in the valuation in some fashion.

4.4.8 Discrete vs. Continuous Rebalancing

Dynamic hedging programs in practice will be rebalanced at discrete intervals or for discrete market moves which is different than theoretical continuous rebalancing and the rebalancing frequency being modelled. In practice, events can happen during hours when markets are closed, resulting in slippage of hedge positions. Consideration would be given to the rebalancing frequency in actual hedging versus modelled hedging and whether this increases or decreases conservatism. The use of a less frequent rebalancing in modelling would tend to create a margin for conservatism within the valuation as there would be larger market moves between rebalancing points in the modelling.

4.4.9 Operational Risk

Due to their complexity, hedging programs can inherently have higher levels of operational risk. Consistent with subsection 2340 of the Standards of Practice, consideration would be given to operational risk when establishing margins for adverse deviation on best estimate assumptions pertaining to the hedging program.

5. PRACTICAL VALUATION CONSIDERATIONS

5.1 Risks Intentionally Not Hedged or Not Modelled, and Hedge Ineffectiveness

There are several mechanical methods that can be employed to include these risks within the valuation. Not all of these methods will work for every valuation methodology. The first four methods do not quantify the risk; they are simply methods to include an amount or margin once quantified. Examples include

- holding a higher CTE level,
- modifying the volatility assumption used,
- modifying the discount rate,
- including an additional cash flow in the valuation, and
- explicitly modelling basis risk in your simulated hedge program payoffs (described in the following section).

5.2 Reflection of Basis Risk

As discussed in section 4.4.1, the common practice of mapping segregated fund asset returns to a linear combination of market index returns, and of similarly mapping hedge instrument returns to a linear combination of market index returns, leads to the result that the model does not provide for basis risk. This shortcoming can be remedied. The most explicit and intuitive way to include basis risk in the model is to add a noise (random) term. The noise term can be added either to the individual segregated fund asset returns (such that they are not perfect linear combinations of the market index returns) or to the hedge instrument returns (such that they are not perfectly aligned with market indices). The magnitude of the noise term can be derived from the regression analysis performed.

For example, in the latter case, the return process for an underlying equity index, which drives simulated hedge program payoffs, can be estimated using regression analysis including an error term. Taking a random draw from the estimated error term of the regression equation, along each step and path, will simulate the underlying equity index return over a time step, thereby adding basis risk to the hedge program payoffs. The simulated hedge program index returns no longer move in a lockstep fashion with the simulated account value returns.

5.3 Determination of Margins for Adverse Deviations

In determining margin for adverse deviations and selecting a methodology to reflect these margins for the various risks and costs associated with a hedging program, the following Standards of Practice are relevant,

> subsection 1740: the precision in the approximation is a consideration in establishing the level of the margin for adverse deviation.
>
> "Selection of a relatively large margin for adverse deviations for the assumption whose uncertainty most affects the calculation and a zero margin for the others may be an appropriate approximation" *(1740.46).*

5.4 Techniques to Decrease Run Time

Stochastic valuation methodologies require significant processing time, especially if a stochastic-on-stochastic method is run, or where scenario testing on stochastic runs is required. Due to the significant processing time, run-time reduction techniques are required in order to make them feasible valuation methodologies. Examples include

- reduce the number of policies modelled with data compression (cluster modelling) techniques,
- reduce the number of scenarios used in both the outer and inner scenario loops, using representative sampling (of scenarios) for real-world valuations and/or random sampling for risk-neutral valuations, and
- lengthen the time-step assumed for rebalancing.

The actuary is reminded that these techniques are approximations and the comments on approximations in section 3 are relevant.

A full discussion on run-time improvements is outside the scope of this Educational Note, but the actuary is encouraged to reference the following documents available on the American Academy of Actuaries' website (www.actuary.org),

- Practice Note on Scenario and Cell Model Reduction, and
- Modeling Efficiency Bibliography for Practicing Actuaries.

5.5 Setting Volatility Assumptions

Many of the approximation methods require a volatility assumption to be determined either as part of the real-world modelling or as part of determination of the risk-neutral liability. The determination of volatility assumptions falls within the mandate of the Working Group on Segregated Fund Calibrations and is out of scope for this Educational Note.

5.6 Use of Dynamic Lapse Functions

Dynamic lapse functions are often employed in real-world modelling of segregated fund investment guarantees. These same functions are often imported into the risk-neutral model in the approximation methods that follow (whether through the investment benchmark or in the approximation method itself as in section 7). For further research on the use of dynamic lapse functions, refer to the paper Modeling and Hedging Dynamic Lapses in Equity-Linked Insurance: A Basic Framework.

5.7 Pros & Cons of Methods Discussed in Sections 6 to 10

Section – Name	Pros	Cons
6 – First-Principles Stochastic-on-Stochastic Method	• Amount of hedges held is explicitly calculated. • Allows for explicit estimation of un-hedged risks.	• Calculation intensive.
7 – Adapted Risk Neutral Method	• Does not require a stochastic-on-stochastic projection. • Does not require development of proxy functions.	• Lack of convergence to CALM when only partial hedging is employed. • Amount of hedges held is not explicitly calculated.
8 – Stochastic-on-Stochastic with Hedge Asset Proxy Method	• No need to explicitly model the assets. • No need to calculate the sensitivities of the investment benchmark (Greeks). • The gains and losses of the hedging strategy are explicitly modelled.	• A thorough understanding of the gains and losses of the hedging strategy and a sufficient history is required to show a stable relationship exists. • Amount of hedges held is not explicitly calculated.
9 – Proxy Function Methods	• No stochastic inner-loop required. • Useful when modelling shorter term guarantees with limited optionality.	• A thorough understanding of the Greeks or gains and losses of the hedging strategy is required. • A sufficient history is required to show a stable relationship exists. • Complicated functions or grids may be required.
10 – Hedge Cost Method	• No stochastic inner-loop required. • Useful when analysing CTE(0).	• Produces a distribution of outcomes that is significantly different from the true outcomes. • Amount of hedges held is not explicitly calculated.

6. FIRST-PRINCIPLES STOCHASTIC-ON-STOCHASTIC METHOD

6.1 Description

The first-principles SOS method is not meant to be an approximation method; it is a first-principles application of CALM when a dynamic hedging program exists, as described in section 4.2. In this method the effect of the dynamic hedging program is calculated using explicit modelling of the hedge positions which are determined from the explicitly-modelled investment benchmark. This is really no different than a fixed cash flow insurance or annuity CALM valuation that attempts to model the asset purchases and sales based on the evolving duration of the liability (i.e., the investment benchmark) along the CALM valuation scenarios.

The investment benchmark can be a real-world or a risk-neutral liability. In either case, one ends up with a requirement for a stochastic-on-stochastic valuation. In one case real-world stochastic (inner loop) is being modelled along real-world stochastic paths (outer loop). In the second, risk-neutral stochastic (inner loop) is being modelled along real-world stochastic paths.

The investment benchmark to be hedged depends on the purpose of the hedging program, but it is common to hedge a measurement of the liability that is based on a risk-neutral or fair value framework and using best estimate assumptions for projecting cash flows. As such, the remainder of this section assumes a risk-neutral inner loop, but the guidance is generally applicable for other contexts as well.

Explicit hedge positions are determined at each node on the real-world paths by determining the sensitivity of the liability to various market moves (i.e., the Greeks), using risk-neutral (stochastic) valuations. The more nodes included in the valuation, the more computationally challenging, but the smaller the time between rebalancing and hence the higher the hedge effectiveness that is modelled. See section 6.3.8 for a discussion on hedge rebalancing.

Having established the hedge positions required at each node, the hedge payoffs at the following time-step are determined by applying the hedge positions to the real-world outer loop. This step is repeated for each node in the real-world outer loop to determine the hedge payoff cash flows, which are then included with the liability cash flows in the CALM valuation.

This method explicitly determines the hedge positions/payoffs and allows for an explicit estimation of unhedged risks and is generally regarded as the first-principles CALM valuation method reflecting hedging. The explicit estimation of the unhedged risk can be quantified by running the model with and without hedging the risk in question.

The major disadvantage of this method is that it is extremely calculation intensive. Consider a relatively realistic (if not overly simplified) example. Assume we need to model the liability cash flows over a 40-year horizon and want to model the rebalancing of the hedge monthly. Also assume we want to hedge equity exposure to three equity markets, and to three points on the yield curve. At each node, we would need a base valuation, one extra valuation for each equity market (for one-sided delta) and one extra valuation for each point on the yield curve (for one-sided key-rate rhos). Assume we determine we need 1,000 outer loop scenarios and 200 inner loop scenarios. The number of projections required for a single insurance contract is then given by 1,000 outer loop scenarios × 40 years × 12 months × 200 inner loop scenarios × 7 assumption sets (base + 3 single equity market shocks + 3 interest rate shocks) = 672,000,000 projections for each insurance contract!

From a practical viewpoint, consideration needs to be given to ways in which the number of calculations can be reduced, e.g., data compression, scenario reductions and increasing the time-step.

6.2 Reflecting Unhedged or Not-explicitly-modelled Risks

Since hedge positions are modelled, the first-principles SOS method lends itself to more explicit modelling of other risks or more explicit inclusion of margin for risk than is the case with some other methods. To the extent that these other risks are material, explicit recognition is preferable but alternatives which may be considered include those available under other approximation methods.

Examples of these implicit methods that are employed under other approximation methods include

- holding liabilities at a higher CTE level,
- modifying the volatility assumption used, and
- modifying the discount rate.

6.3 Risks and Costs in a Hedging Program to Reflect in the Valuation

The following section expands upon section 4.4 for the first-principles SOS method.

6.3.1 Basis Risk

In addition to performing fund mapping studies to analyse basis risk, the actuary is also able explicitly to model, test and analyse basis risk under this methodology.

Basis risk can be reflected in the valuation by explicitly including an additional index within the hedge payoff. For example, assume that the product's fund mix is mapped to 50% TSX, 30% S&P 500 and 20% EAFE, and hence liability cash flows would vary along the real-world paths based on these indices. Then, each of the following would simulate the inclusion of basis risk in the valuation,

- adding a term to the hedge payoff that is unrelated to these indices, (For example, model hedge payoffs based on mix of 45% TSX, 30% S&P 500, 20% EAFE and 5% in an unrelated (noise) index.)
- adding a noise term to the liability fund mix instead of the hedge payoff mix, and
- increasing the volatility of one or more of the existing indices, but only for purposes of modelling either the liabilities or the hedge payoffs.

6.3.2 Liquidity Risk, Bid-Ask Spread and Market Effect Costs

Since hedge positions are modelled, it is possible to apply a margin to bid-ask spread assumptions and use this in conjunction with the modelled trading volumes to capture these costs. It may be appropriate to vary the margin based on current environment in the real-world path to reflect potential difficulty of trading. The costs included would be dependent on hedging instruments utilized.

6.3.3 Transaction Costs and Commissions

As mentioned previously, hedging programs can have significant amounts of trading. This methodology allows for explicit reflection of trading costs as the trading volumes are calculated within the modelling. Adjustments may be necessary if modelled rebalancing frequency does not match actual rebalancing frequency.

6.3.4 Counterparty Risk

This risk can be included in the valuation in a similar fashion to an insurance or fixed annuity CALM valuation by applying a margin to hedge payoffs where appropriate.

6.3.5 Volatility

It may be appropriate to link future implied volatilities to the environment on the real-world path (volatility may vary along real-world paths if different volatility regimes are possible or if stochastic volatility is used) if practical where options are employed.

6.3.6 Risks Intentionally Not Hedged

Practically speaking, the target hedge may not equal the full risk-neutral liability as portions of the risk may not be hedged. It may, therefore, be appropriate to apply a factor to the calculated Greeks or some other modification to reflect these unhedged risks. For example, if only 75% of the interest rate exposure is hedged as per the investment policy, it would be straightforward to project the hedge portfolio to match only 75% of projected rhos.

In addition, there may be CALM liability cash flows that are not included in the investment benchmark (e.g., PfADs not included in the risk-neutral liability or a hedge policy that ignores the hedging of fees and only hedges benefits). Net cash flows along real-world paths will reflect the effects of any risks omitted from the hedging strategy since the liability cash flows will encompass all aspects of the liability, while the hedge payoffs will only cover off the hedged items (captured by the hedge target).

6.3.7 Risks Not Explicitly Modelled

This first-principles SOS method allows for more explicit modelling of risks than many other methods. There may still be risks that are intentionally not modelled in order to simplify or improve run times. These risks would still need to be included in the valuation. The methodologies described in section 5.1 may be appropriate for reflecting these risks.

6.3.8 Discrete vs. Continuous Rebalancing

The reality of not being able to rebalance continuously adds risk to the hedging program versus theoretical continuous rebalancing. However, when modelling the hedging program, the rebalancing assumed/modelled is often less frequent than actually done in practice. If this is the case, an additional implicit margin is introduced into the valuation. The additional implicit margin is introduced because the less frequent rebalancing results in larger losses as larger market moves would be expected in the longer elapse time between nodes. The larger losses occur because the hedges have a more linear change in value than does the liability (gamma is smaller for the hedge portfolio than the liabilities). Thus, the first-principles valuation method would be expected to generate a higher reserve with less frequent rebalancing modelled (i.e., fewer nodes).

6.4 Other Considerations

6.4.1 Modelling the Risk-neutral Liabilities along the Real-world Paths

The real-world environment at each node along the real-world path is an important piece of information to reflect in the calculations required at that node. The real-world market and interest rate levels are used in determining the liability cash flows as well as the policyholder behaviour. The real-world equity volatility and interest rate levels can be used as inputs to determine the risk-neutral liability. Modifications may be required to achieve good alignment between the risk-neutral liability that is used in practice for hedging and the risk-neutral liability that is being modelled along the real-world path in CALM.

7. ADAPTED RISK-NEUTRAL METHOD

7.1 Description

In the context of an approximation method to CALM, a pure risk-neutral approximation method is not appropriate since this liability would not be dependent on the assets and reinvestment strategy backing it, unless the investment strategy was to purchase long-dated options at market cost that replicated the liabilities closely. In this document we will refer to an "adapted" risk-neutral method. The term "adapted" is used to represent the fact that modifications from a pure risk-neutral methodology are required in order to approximate CALM. The CALM liability is dependent on the hedging strategy employed whereas a pure risk-neutral liability would not be dependent on the hedging employed. This means that the CALM liability will converge towards a risk-neutral liability as more and more aspects of the liability are hedged. The adapted risk-neutral approximation method can thus be an appropriate approximation when material hedging is performed against a risk-neutral liability.

In theory, if all aspects of the risk-neutral liability are being hedged (all Greeks of the investment benchmark are hedged), the result of running a first-principles SOS valuation would be similar to the risk-neutral liability. This is because the change in the risk-neutral liability is what is experienced along every real-world path in the first-principles SOS valuation. When all aspects of the risks are hedged, the present value of each of the real-world scenarios will converge towards the risk-neutral liability. Convergence between first-principles SOS and the risk-neutral liability is more likely if there is good consistency between the real-world outer loop and the risk-neutral inner loop (e.g., similar assumptions) and all aspects of that risk-neutral liability are hedged (benefits and fees, linear and non-linear risks, and volatility, equities and interest rates).

Adaptations may be required if aspects of the risk-neutral liability are not hedged (e.g., if fees are not hedged, valuing the fees separately in a real-world CALM valuation may be appropriate). Adaptations may also be required to reflect imperfect policyholder behaviour. Other adaptations that may be required from a pure risk-neutral approach could include use of a discount rate that exceeds the risk-free interest. Margins would be required on these real-world assumptions.

Underlying the determination of a risk-neutral liability is an assumption with respect to equity volatility as well as interest rates and interest rate volatility. Establishing these assumptions is beyond the scope of this Educational Note. Under a CALM valuation, the hedging strategy employed does affect the liability and thus it would be appropriate that companies using different hedge vehicles could come up with a different risk-neutral liability. Where options form a large portion of the hedging program implied volatility grows in importance relative to realized volatility in establishing a volatility parameter.

One major advantage of the adapted risk-neutral methodology is it does not require a stochastic-on-stochastic projection. It also provides good alignment between asset and the liability movements and could reduce income volatility.

The major disadvantage of the adapted risk-neutral methodology is the lack of convergence to CALM when only partial hedging is employed. Another disadvantage is that the amount of hedges held is not explicitly calculated when using the risk-neutral approximation method. This limits the methods by which one can capture and quantify other risks that might be related to trading frequency or amount of assets or hedge contracts held.

7.2 Reflecting Unhedged or Not-explicitly-modelled Risks

Unlike other methods, holding a CTE level higher than CTE(0) has no real meaning in a risk-neutral valuation. As such, other methods are required in order to reflect unhedged risks or risks and costs that have not been considered explicitly in the reserve. Margins for adverse deviation would also be added to assumptions that are real-world in nature such as when spread is added to a risk-free discount rate.

Examples of methods that could be used to reflect unhedged risks are

- modifying the volatility assumption used,
- modifying the discount rate, and
- including an additional cash flow in the valuation.

7.3 Risks and Costs in a Hedging Program to Reflect in the Valuation

The following section expands upon section 4.4 for the adapted risk-neutral method.

7.3.1 Basis Risk

Implementing basis risk into the adapted risk-neutral approximation method is not as intuitive as other methods since the hedge assets and/or payoffs are not explicitly modelled. As such, once a level of margin is determined for basis risk through fund mapping or other analysis, there are several ways to include it in the valuation. These methods are described in section 7.2.

7.3.2 Liquidity Risk, Bid-Ask Spread and Market Effect Costs

Under this method, there is no explicit modelling of transactions and, as such, these costs cannot be calculated explicitly within the valuation. These costs may be included using methods described in section 7.2.

7.3.3 Transaction Costs and Commissions

Bid-ask spread, transaction costs and commissions are expenditures required when buying or selling assets to rebalance a hedge portfolio. In the adapted risk-neutral approximation method, the cost of these transactions cannot be directly calculated in the modelling. These costs may be included using methods described in section 7.2.

7.3.4 Counterparty Risk

Again, due to the lack of explicit modelling of transactions, these costs cannot be explicitly modelled in the valuation. These costs may be included using methods described in section 7.2.

7.3.5 Volatility

Under CALM, the hedging or investment strategy affects the liability value. Where the hedging program utilizes options, CALM would reflect the cost of these options in the liability value. Said another way, the CALM liability would be dependent upon the implied volatility of these options. Where the dynamic hedging program does not utilize options, CALM would not be dependent on the implied volatility and only the realized volatility. Thus, when establishing the volatility parameter within the adapted risk-neutral valuation consideration would be given to the hedging strategy that is to be employed. Where options are employed the liability would be more heavily dependent on implied volatility than in a situation where options do not form a material part of the hedging program. This is clearly a framework different from a pure market consistent or risk-neutral framework as described in the 2010 task force report.

7.3.6 Risks Intentionally not Hedged

A hedge bifurcation may be an appropriate option under the adapted risk-neutral method. In the case of unhedged fees, bifurcating the liability and calculating the hedged benefits on a risk-neutral basis and the unhedged fees on a real-world basis may be appropriate. Care would be taken in this type of bifurcation to avoid disconnecting insurance contract feature interactions. Where bifurcation is performed, some of the above risks that vary based on scenario could be captured within the real-world bifurcated portion of the liability. Generally, care would be taken when modelling risks using a combination of real-world and risk-neutral environments.

7.3.7 Risks Not Explicitly Modelled

Many of the risks described above are not explicitly modelled under this method. There may be other risks not listed above and these may be included within the valuation using techniques described in section 7.2.

7.3.8 Discrete vs. Continuous Rebalancing Frequency

The adapted risk-neutral method would not naturally capture the slippage caused by less than continuous rebalancing. Also, this method would not include any conservatism in the valuation by modelling a less frequent rebalancing than actually used in practice. As such, the actuary would consider including an additional cost in the valuation using one of the methods described in section 7.2.

7.4 Other Considerations

7.4.1 Use of Dynamic Lapse Functions

Dynamic lapse functions are often employed in real-world modelling of segregated fund investment guarantees. These same functions are often imported into the adapted risk-neutral approximation method. For further research on the use of dynamic lapse functions, refer to the paper, Modeling and Hedging Dynamic Lapses in Equity-Linked Insurance: A Basic Framework.

8. STOCHASTIC-ON-STOCHASTIC WITH HEDGE ASSET PROXY METHOD

8.1 Description

This method is a simplification of the first-principles stochastic-on-stochastic method, which uses a relation to express the market value of the hedge portfolio (target hedge) as a function of the investment benchmark. Similar to the first-principles method, the investment benchmark could be a real-world or a risk-neutral liability. With this method, the hedge portfolio (target hedge) is defined entirely by the investment benchmark, eliminating the need to determine and then model the hedge portfolio specifically. Thus, as you progress from one node to the next the hedge gains and losses are determined by applying the relation to the change in the investment benchmark. This eliminates the need to perform multiple shocked valuations of the investment benchmark that are required under the first-principles methodology in order to calculate the Greeks that are necessary to determine the composition of the hedge portfolio.

With the first-principles SOS method, the sensitivities of the liability (i.e., the Greeks) directly determine the composition of the hedge portfolio and the transactions carried out to rebalance the portfolio. Since this approximation method does not model the hedge assets themselves, it is no longer necessary to calculate the sensitivities of the investment benchmark. However, the SOS with hedge asset proxy method is limited in its ability to model the risks related to rebalancing the hedge portfolio, including the risk of a widened spread and counterparty risk.

The logic behind this method is that the market value of the hedge portfolio should be expected to fluctuate in a random but comprehensive manner around the investment benchmark. Thus, the gains and losses of the hedge portfolio can be calculated using some function of i) the change in the investment benchmark, and ii) the changes in the economic environment on the outer loop. These gains and losses are then added to the liability valuation cash flows along each real-world path.

A number of relations or functions can be used to model the market value of a hedge portfolio, but the two described below have the advantage of being intuitive.

- The market value of the hedge portfolio is equal to the investment benchmark, to which we add a semi-random error component corresponding to anticipated behaviour of the gains and losses of the hedging strategy. The semi-random component might differ from a white noise distribution to the extent that one can derive a relationship between these hedge errors and the market conditions (e.g., size of equity or interest rate moves since last node).
- The change in market value of the hedge portfolio is equal to a proportion of the variation of the investment benchmark.

A more complex relation based on different market indicators might also be used if it improves the quality of the valuation.

The SOS with hedge asset proxy method provides the advantages that

- there is no need to explicitly model the assets,
- there is no need to calculate the sensitivities of the investment benchmark (Greeks), and
- the gains and losses of the hedging strategy are explicitly modelled.

On the other hand, this method requires a thorough understanding of the gains and losses of the hedging strategy being analyzed in order to ensure that the strategy is properly modelled in the valuation. This method can be considered a valid approximation for the simulation of a hedging strategy with a sufficient history that shows that a stable relationship exists. However, the actuary would be careful when using historical data to calibrate this method since the historical data can reflect a particular economic reality that is not suitable for projection. The actuary would also be careful when applying this approximation method to a new hedging strategy or covering a new product with a risk profile different from those that are known. The actuary is referred to section 3 on the appropriateness of approximation methods to CALM. Also, the relationship of the hedge gains and losses to the investment benchmark may change over time and with changing market conditions. As such, care woud be taken to ensure the relationship does not break down over time and across scenarios being modelled.

8.2 Reflecting Unhedged or Not-explicitly-modelled Risks

To model unhedged or not-explicitly-modelled risks, the actuary may choose to adjust the parameters of the relation to increase the potential losses of the hedging strategy. Otherwise, the methods available with the first-principles SOS method also apply.

8.3 Risks and Costs in a Hedging Program to Reflect in the Valuation

The following section expands upon section 4.4 for the SOS with hedge proxy asset method.

8.3.1 Basis Risk

Once quantified, these risks can be explicitly modelled with the SOS with hedge proxy method using the relation that models the market value of the assets. This relation is an integral part of the method.

There might be an insufficient margin on mortality and lapse assumptions if this method is based on historical data that is representative of best estimate experience. In that case, an additional margin would be added to the valuation. If the model has been calibrated with the first-principles SOS method, this additional margin might not be necessary as it may already be captured.

8.3.2 Liquidity Risk, Bid-Ask spread and Market Effect Costs

Since hedge positions are not modelled explicitly under this method, it is not possible to use the modelled trading volumes to capture these costs. It may be appropriate to modify the function that is applied to the investment benchmark to capture these costs or use methods as described in section 6.2.

8.3.3 Transaction Costs and Commissions

Transaction costs and commissions are expenditures required when buying or selling assets to rebalance a hedge portfolio. Under this approximation method, transactions are not modelled so this approach cannot be used.

One way of including transaction costs is to project a fixed or variable cost in the valuation reflecting as closely as possible the experience of the insurer with regard to these transaction costs as a whole. Two of the methods available to the actuary to model these costs are

- setting a fixed cost per given period, and
- setting a variable cost subject to market volatility. Market volatility influences transaction costs in two ways; 1) more transactions are made in order to rebalance the hedge fund portfolio, and 2) the bid-ask spread may widen.

8.3.4 Counterparty Risk

This risk cannot be explicitly modelled with this approximation method because the hedge portfolio itself is not modelled. The actuary would add a provision to cover this risk (see section 4.4.4).

8.3.5 Volatility

If using either a real-world or risk-neutral investment benchmark, it may be appropriate to link future volatilities used in calculating the investment benchmark to the environment on the real-world path (volatility may vary along real-world paths if different volatility regimes are possible or if stochastic volatility is used).

8.3.6 Risks Intentionally Not Hedged

In practice, the target hedge may not include all Greeks of the risk-neutral liability or real-world liability as portions of the risk may not be hedged. It may, therefore, be appropriate to define the target hedge as a function of the risk-neutral or real-world liability but not equal to it.

8.3.7 Risks Not Explicitly Modelled

Given that this methodology does not explicitly calculate the hedge positions, there may be risks that are not explicitly modelled. These risks would still need to be included in the valuation. The methodologies described in section 5.1 may be appropriate for reflecting these risks.

8.3.8 Discrete vs. Continuous Rebalancing

Since the gains and losses of the hedging strategy are directly modelled using a relationship, it is no longer necessary to link the rebalancing frequency with the projection period in the valuation. For example, the quarterly gains and losses of a weekly rebalanced hedging strategy can be analyzed and modelled. This can help make calculations much faster.

We encourage the actuary to analyze the gain and loss profile on different time horizons (daily, weekly, monthly, and quarterly) to have a thorough understanding of the underlying volatility of the gains and losses of the hedging strategy before choosing a time step in the valuation.

9. PROXY FUNCTION METHODS

9.1 Description

There are a number of approximation methods which involve replacing the stochastic inner loops (the risk-neutral calculations) required in a stochastic-on-stochastic method. Different variations replace the inner loops with closed form functions or pre-calculated grids used to approximate the investment benchmark(s) and/or asset sensitivities. This is going one step further down the approximation path than the method described in section 8 (SOS with hedge asset proxy) since in this case there is no stochastic inner loop calculation while in section 8 the inner loop still remains.

Proxy function approximation methods include those using one or more closed form functions to provide for measurement of the investment benchmark sensitivities at future dates along the real-world outer loop paths. Multiple proxy functions may be used to approximate better the results of the inner risk-neutral loop in different environments, or at differing time horizons along the real-world paths. The closed form function(s) would be used to derive all necessary sensitivities of the investment benchmark(s), i.e., the Greeks, in order explicitly to determine the required volumes of hedge instruments necessary at future dates along the real-world paths to meet hedge strategy tolerances. By determining explicit hedge positions at each node on the paths, hedge payoffs are readily available.

Another variation of this method is the use of a grid of the required hedge positions under a range of environments that may develop along the real-world outer loop paths. The grid, which may include a number of dimensions to capture sufficiently the effect of the many factors that can affect investment benchmark sensitivities such as time to maturity, degree of being “in-the-money”, changing volatility or interest rate level (if applicable), etc. The grid would replace the on-the-fly calculation of required hedge positions from closed form functions.

The addition of a proxy for the hedge asset payoff to the liability proxy function approach has been suggested as a simplification that eliminates the need for calculation of specific hedge asset positions. If the hedge positions are desired, a proxy for the investment benchmark is used to calculate Greeks and from this proxy the hedge portfolio is determined. If the hedge positions are not desired, then a proxy for the hedge payoffs is used and no hedge positions or Greeks are explicitly calculated.

The proxy function approximation method is often considered in connection with the use of replicating portfolios. These replicating portfolios would consist of a collection of instruments (some of which may be only hypothetical and not actively traded in financial markets) whose sensitivities are expected to reproduce those of the investment benchmark(s). Functions based on replicating portfolios might be expected to include linear, quadratic/exponential and optional components in order to capture the hedging results in a widest range of possible future real-world environments.

9.2 Reflecting Unhedged or Not-explicitly-modelled Risks

Since hedge positions are modelled, except in variants where hedge asset proxies are also included, this method lends itself to more explicit modelling of other risks than is the case with some other approximation methods. The comments and methods included in section 6.2 and, in cases where a hedge asset proxy is included, in section 8.2 are therefore applicable.

To the extent that these other risks are material, explicit recognition is preferable but alternatives which may be considered include those available under other approximation methods.

9.3 Risks and Costs in a Hedging Program to reflect in the Valuation

The risks and costs to be reflected when using a proxy function method are the same as those inherent in the first-principles SOS methodology. While the implementations of this method that track hedge assets explicitly use different methods in determining those assets from the first-principles SOS method, the volumes of those hedge assets are nonetheless available. As such, the techniques for reflecting the various risks and costs used in a first-principles SOS method can also generally be applied when a proxy function is used for the investment benchmark(s) and the comments of section 6.3 are applicable. Implementations of the proxy function method which do not track hedge asset trading explicitly, such as those including a hedge asset proxy as well, will require alternate techniques to provide for some risks. The comments of section 8.3 may be relevant in these circumstances.

9.4 Other Considerations

While the theoretical appeal of approximating the inner stochastic loops with either closed form function(s) or grids is evident, the practical difficulties involved in using such approximation methods in the context of the multi-faceted guarantees offered in most segregated fund products are significant. With the many policyholder options, economic factors and long time-frames often involved, the necessary functions or grids are likely to be quite complex. Extensive testing would be required to confirm that these types of approximation methods are sufficiently robust to respond realistically in the wide range of potential environments represented by the outer real-world stochastic loops.

It may be more practical to use these approximation methods in the context of shorter-term guarantees with limited optionality. A maturity guarantee with a short remaining term and no policyholder reset option might, for example, be a potential candidate for this treatment.

Apart from these unique aspects of this family of approximation methods, many of the considerations present in the methods discussed in other sections of this educational note are also applicable.

10. HEDGE COST METHOD

10.1 Description

The hedge cost method has a serious drawback. Under the hedge cost method, the real-world scenarios resulting in adverse outcomes will be those of poor investment returns, similar to a situation where there is no hedging in place. More detail on this drawback is described in section 10.4.

The hedge cost method is an approximation method to CALM that reflects the effects of a hedging program in two ways, the cost of the hedging and the benefits from the hedging program. Instead of using a stochastic-on-stochastic (SOS) approach, this method uses simplifying assumptions and one set of real-world stochastic scenarios to reflect the dynamics of a hedging program. The following is the definition of the method taken from the 2010 task force report (section 5.5.2).

> This method uses stochastic methodology with two simplifying assumptions,
>
> > it recognizes the cost of hedging in the form of basis points of expense, as opposed to dynamically estimating the hedge costs using nested risk-neutral stochastic paths, and
> >
> > it recognizes the benefits of hedging in the form of a percentage of all future guarantee top-ups that will be offset by future hedging gains. This percentage is a measure of the expected effectiveness of the hedging program.
>
> The above two assumptions would be developed as all other assumptions, drawing on the best available information and expertise. These assumptions need not be scalars — they can vary by duration or other attributes as appropriate. Each of these assumptions would have a MfAD appropriately recognizing the potential for misestimation of the best estimate or deterioration of the best estimate assumption.

The cost of the hedging program is typically expressed as a basis point measure. All of the costs of the hedging program would be considered. Due to the diverse risks within a hedged segregated fund product and the hedging program itself, the basis point measure for each risk may not have the same base (e.g., account value versus guarantee benefit base). The base should behave in a similar way as the risk that is being addressed.

The benefit of the hedging program is a decrease in the expected costs of the guarantee top-ups. This can be considered as the hedge effectiveness of the program, and would account for items that are not already reflected in the cost of the hedging program. The decrease in the expected costs is typically expressed as a percentage reduction in the expected guarantee top-ups.

The amount of the basis point cost and the percentage benefit may also vary over the duration of the projection. These values are not required to be static scalars, since the costs and benefits may not be constant over the projection period.

If the pattern of the hedge cost liability does not exactly match the pattern of the hedged liability, the basis points cost may need to be re-calibrated in the future. This would reflect the observed experience and current environment. The percentage of expected hedge benefit may also change over time as experience from the hedging program is realized.

Even with these adjustments, the hedge cost method produces a distribution of outcomes that is significantly different from the true outcomes and the emergence of profit and tracking of hedge error is not expected to follow their actual patterns. The actuary would be cautious when using this method over a long period of time.

The hedge cost method has a benefit of reflecting the risks of a hedging program in a computationally efficient manner. This is due to only one set of scenarios being projected and not requiring stochastic-on-stochastic calculations. The computational efficiency is at the cost of a more rigorous risk calculation.

10.2 Reflecting Unhedged or Not-explicitly-modelled Risks

Due to the use of a real-world scenario projection and guarantee top-ups being reflected in this method (although at a lower amount), a higher CTE level will produce a higher liability under the hedge cost method. The CTE level that is used would consider, and be consistent with, the basis point measure of the costs and the percentage used in the calculation of the hedging benefit. Due to the nature of the hedge cost method and the fact that is does not reflect the true distribution of results, the CTE level will largely be driven by poor returns which may not be the driver of the actual costs of the hedging portfolio.

10.3 Risks and Costs in a Hedging Program to Reflect in the Valuation

The following section expands upon section 4.4 for the hedge cost method.

10.3.1 Basis Risk

Basis risk can be based on a best estimate assumption based on historical experience from the hedging program and/or correlation analysis between the actual fund performance and the performance of the indices to which it is mapped (i.e., the hedge assets). An appropriate margin for adverse deviation can also be determined from this information. This margin would reflect the basis risk that may be experienced under severe conditions.

Basis risk is typically reflected as a basis point cost that is a function of the account value.

10.3.2 Liquidity Risk, Bid-Ask Spread and Market Effect Costs

Since hedge positions are not modelled explicitly under this method, it is not possible to use the modelled trading volumes to capture these costs.

10.3.3 Transaction Costs and Commissions

Transaction costs and commissions are expenditures required when buying or selling assets to rebalance a hedge portfolio. These costs may be directly calculated when the asset transactions are explicitly modelled. Under the hedge cost method, asset transactions are not explicitly modelled so bid-ask spread, transaction costs and commissions cannot be calculated.

One way of getting around this problem is to project a fixed or variable cost in the valuation reflecting as closely as possible the experience of the insurer with regard to these transaction costs as a whole. Two potential ways to model these costs are

- setting a fixed cost per given period, and setting a variable cost subject to market volatility. Market volatility influences transaction costs in two ways, 1) more transactions are made in order to rebalance the hedge fund portfolio, and 2) the spread may widen.

These costs can be modelled as a basis point cost or as a reduction in the assumed hedge effectiveness (i.e., reduction in expected claims).

10.3.4 Counterparty Risk

Consider including transaction costs and commissions in a fashion similar to bid-ask spread (liquidity risk).

10.3.5 Volatility

The risk-neutral and real-world volatility assumptions need both to be considered in that

the volatility assumed in the risk-neutral projection determines the basis point hedge cost that is used in the real-world projection, and

the volatility assumed in the real-world projection determines the level of claims prior to the hedge payoff being applied.

The interactions between these two volatility assumptions would be tested and understood.

10.3.6 Risks Intentionally Not Hedged

The real-world projection would implicitly (e.g., policyholder behaviour margins would be included even if only best estimate policyholder behaviour is assumed in the investment benchmark) or explicitly (e.g., additional costs or higher CTE level) model any risks that are intentionally not hedged.

10.3.7 Risks Not Explicitly Modelled

Given that this methodology does not explicitly calculate the hedge positions, there may be risks that are not explicitly modelled. These risks would still need to be included in the valuation. The methodologies described in section 5.1 may be appropriate for reflecting these risks.

10.3.8 Discrete vs. Continuous Rebalancing

The hedge cost method does not project the rebalancing of hedge portfolio. Therefore, this method does not capture any conservatism in the valuation by modelling a less frequent rebalancing than is actually used in practice. As such, one would consider including an additional cost in the valuation using one of the methods described in section 5.1.

10.4 Other Considerations

This method has a serious drawback. Under the hedge cost method, the real-world scenarios resulting in adverse outcomes will be those of poor investment returns, similar to a situation where there is no hedging in place. This is because the hedging imperfection under this method is allowed for by assuming that a fraction of the guarantee benefits is not covered by the hedging. This technique, coupled with the CTE measure, is likely to result in insurance contract liabilities that are larger than the best estimate cost of hedging, i.e., it includes a PfAD. However, under a typical delta hedging strategy, scenarios resulting in adverse outcomes are not necessarily those of poor investment returns. Adverse scenarios under a delta hedging strategy are those where large market shifts, either upward or downward, occur at times where the gamma of the guarantee liability is large, i.e., where the liability function exhibits a large convexity. Therefore, there is no assurance that the PfAD provided by the hedge cost method is proportionate to the risk for which a provision should be established. The working group is of the opinion that this is a fundamental flaw and that the hedge cost method would be a method companies transition away from for establishing reserves. For analysis using CTE(0), this method is not flawed and as such could be useful for planning, pricing or other analysis.

11. BIBLIOGRAPHY

Modeling Efficiency Work Group of the American Academy of Actuaries. Scenario and Cell Model Reduction. 2010.

———. Modeling Efficiency Bibliography for Practicing Actuaries. 2010.

Panneton, Christian-Marc, and Boudreault, Mathieu. "Modeling and Hedging Dynamic Lapses in Equity-Linked Insurance: A Basic Framework". *Risks & Rewards* 58 (August 2011): 1–10.

12. APPENDIX: NUMERICAL EXAMPLES

The following examples are for purposes of aiding the understanding of the description of the methodologies. They are not meant to be representative or indicative of any method.

12.1 Description

The liability used in the examples that follow is a five-year at the money put option with fees based on the guarantee value. The risk-neutral liability for claims has been calculated using the Black-Scholes formula. An equity delta has been calculated as well as an interest rate rho.

	path \ time	0	1	2	3	4	5
A	RW path - Equity Return		8.0%	8.0%	-10.0%	-15.0%	8.0%
B	RW path - Interest Rate Level	4.0%	4.0%	4.0%	3.5%	3.5%	3.5%
C	MV	1,000	1,080	1,166	1,050	892	964
D	GV	1,000	1,000	1,000	1,000	1,000	1,000
E	Liability cash flows (fee portion) (based on GV)	1.0%	10.0	10.0	10.0	10.0	10.0
F=(D-C)	Liability cash flows (claims portion)						36.3
G	RN Liability (claims portion)	87.1	65.8	43.2	62.8	116.7	
H	+1% equity	(2.5)	(2.3)	(1.9)	(3.0)	(5.4)	
I	-1% equity	2.6	2.4	2.0	3.1	5.6	
J=(H-I)/2	Delta	(2.5)	(2.3)	(2.0)	(3.0)	(5.5)	
K	+10 bps	(1.6)	(1.1)	(0.7)	(0.7)	(0.6)	
L	- 10 bps	1.7	1.2	0.7	0.7	0.6	
M=(K-L)/2	Rho	(1.6)	(1.2)	(0.7)	(0.7)	(0.6)	
	Black-Scholes parameters used to create G:						
	Effective annual interest rate	4.00%	4.00%	4.00%	3.50%	3.50%	
	Volatility	20.0%	20.0%	20.0%	20.0%	20.0%	
	Term	5	4	3	2	1	

12.2 Example: First-principles SOS Method

The following shows the projected liability and hedge payoffs for the sample real-world path shown in section 12.1. For simplicity, interest earnings on the hedge positions have been ignored and only gains due to market movements have been shown.

The investment benchmark in this first example attempts to hedge 100% of the delta and rho of the risk-neutral liability.

	cash flow \ time	**PV**	1	2	3	4	5
E	Liability cash flows (fee portion) (based on GV)	44.5	10.0	10.0	10.0	10.0	10.0
F	Liability cash flows (claims portion)	29.9	-	-	-	-	36.3
=JxA/1%	Hedge payoffs - short futures (equity change impact only)	65.0	(18.7)	(15.8)	30.3	82.6	-
=MxB/1%	Hedge payoffs - swaps (interest rate change impact only)	3.1	-	-	3.5	-	-

The investment benchmark in this next example attempts to hedge 75% of the delta and 50% of the rho of the risk-neutral liability.

	cash flow \ time	**PV**	1	2	3	4	5
K	Liability cash flows (fee portion) (based on GV)	44.5	10.0	10.0	10.0	10.0	10.0
F	Liability cash flows (claims portion)	29.9	-	-	-	-	36.3
=JxAx0.75/1%	Hedge payoffs - short futures (equity change impact only)	48.7	(14.0)	(11.8)	22.8	61.9	-
=MxB*0.50/1%	Hedge payoffs - swaps (interest rate change impact only)	1.6	-	-	1.8	-	-

This is only one example of a scenario path that simply lays out the idea behind what is being modelled under each of the many scenarios being run under this method.

12.3 Example: Adapted Risk-neutral Approximation Method

In the adapted risk-neutral example, the investment benchmark is the risk-neutral liability for claims. Thus, the best estimate liability is estimated as the risk-neutral claim liability less the real-world fee liability. Margins are not explicitly calculated in this method, so would need to be added by other means.

		PV	1	2	3	4	5
G	Best Estimate Reserve (risk neutral) - Claims	87.1					
E	Best Estimate Reserve (real world) - Fees	44.5	10.0	10.0	10.0	10.0	10.0
G-E	Best Estimate Reserve - Total	42.6					
	MfADs	X					

12.4 Example: SOS with Hedge Asset Proxy Approximation Method

The following shows the projected liability and hedge payoffs for the sample real-world path shown in section 12.1. For simplicity, interest earnings on the hedge positions have been ignored and only gains due to market movements have been shown.

Here again, the investment benchmark is the risk-neutral liability for claims. The proxy for the hedge gains and losses is 90% of the change in the risk-neutral liability less 2: [90% x (RNt – RNt-1) – 2]. This particular proxy has no meaning, and is much simpler than a proxy that would be used in practice, and is only shown to understand the general concept. Note that this approximation method requires the calculation of a risk-neutral liability, but not the corresponding Greeks.

	cash flow \ time	PV	1	2	3	4	5
E	Liability cash flows (fee portion) (based on GV)	44.5	10.0	10.0	10.0	10.0	10.0
F	Liability cash flows (claims portion)	29.9	-	-	-	-	36.3
$=0.90x(G_t-G_{t-1})-2$	Hedge payoff	12.7	(21.2)	(22.3)	15.6	46.5	

12.5 Example: Proxy Function Approximation Method

The following shows the projected liability and hedge payoffs for the sample real-world path shown in section 12.1. For simplicity, interest earnings on the hedge positions have been ignored and only gains due to market movements have been shown.

Here again, the investment benchmark is the risk-neutral liability for claims. The proxy for the hedge gains and losses is two times the equity return less seven times the interest rate change. [2 x equity return – 7 x interest rate change]. This particular proxy has no meaning, and is much simpler than a proxy that would be used in practice, and is only shown to understand the general concept. Note that this proxy does not require the calculation of a risk-neutral liability. Also note that this is an example of using an asset hedge proxy where the Greeks are not explicitly calculated at each node.

	cash flow \ time	PV	1	2	3	4	5
E	Liability cash flows (fee portion) (based on GV)	44.5	10.0	10.0	10.0	10.0	10.0
F	Liability cash flows (claims portion)	29.9	-	-	-	-	36.3
$=(2xA-6x(B_t-B_{t-1}))/1\%$	Hedge payoff	2.8	(16.0)	(16.0)	23.0	30.0	(16.0)

12.6 Example: Hedge Cost Approximation Method

Here again, the investment benchmark is the risk-neutral liability for claims. The liability is calculated as the PV of the hedge costs, plus the claims cash flows less the fees and hedge payoff. The hedge payoff here is assumed to be 75% of the claims paid. Note that in practice the hedge costs would typically be calculated as a bps charge on the account value and projected out over the scenario path and discounted back to the valuation date. For simplicity the PV of the risk-neutral value of the hedges has been used.

		PV					
G	Hedge Costs	87.1	-				
E	Liability cash flows (fee portion) (based on GV)	44.5	10.0	10.0	10.0	10.0	10.0
F=(D-C)	Liability cash flows (claims portion)	34.9					36.3
=0.75xF	Hedge payoff	26.2					27.2
=G+F-0.75xF-E	Reserve	51.3					