Insurance and Actuarial Advisory Services

Market Value Margins for Insurance Liabilities in Financial Reporting and Solvency Applications

October 1, 2007

Ernst & Young

Quality In Everything We Do

Table of Contents

Executive Summary

Introduction

Two important insurance industry developments have brought considerable attention to the need for a consistent framework for measuring the fair value of insurance liabilities. First, new financial reporting standards aimed at producing *fair value* financial statements for insurers are being considered by the IASB, the FASB and other accounting policy organizations. Second, insurer solvency standards being considered in various countries have focused on the potential variability of fair value financial statements as the basis for establishing minimum regulatory capital requirements for insurers.

A key element of the calculation of fair values for insurance liabilities is the determination of a market value margin (MVM) for non-hedgeable risk. The *cost of capital method* has emerged as a strongly preferred methodology for this calculation and has already been adopted in certain jurisdictions as part of the framework for establishing capital requirements and measuring available capital.

This report examines the cost of capital method in two specific contexts:

- Measuring the fair value of insurance liabilities for financial reporting purposes, and
- Establishing regulatory capital standards for insurers.

Examples of the application of the cost of capital method to life and non-life insurance products are used to highlight practical challenges that must be addressed to ensure proper implementation of this method.

Overview of the Cost of Capital Method

The cost of capital method defines the market value margin (MVM) as the present value sum of a capital base each period, over the lifetime of the liabilities, multiplied by a cost of capital rate per period. The inputs required for the calculation include:

- Initial Capital Base
- Capital Base in Each Subsequent Period Over the Lifetime of the Liability
- Cost of Capital Rate Per Period

The cost of capital approach to calculating the MVM is conceptually appealing and intuitive. However, as summarized below, for the financial reporting and solvency applications addressed in this report it is important to apply the methodology consistently, to validate any approximations and simplifications that are adopted, and to calibrate key assumptions, such as the combined effect of the capital base and the cost of capital rate, so that the resulting fair value estimates reasonably match observed benchmarks for the price of insurance risk.

Major Findings

The following briefly summarizes the major findings contained within the report:

- <u>*Key Parameters Used in the Cost of Capital Methodology and in Pricing Practices Should Be Reconcilable*</u>

  The cost of capital method closely parallels the methods that many insurers currently use to price insurance and reinsurance products and business transactions, which is one reason why this method is advocated as the basis for calculating the MVM for non-hedgeable insurance risk. We recognize that other adjustments to the MVM are likely necessary to obtain a value that is consistent with market prices for insurance liabilities. Nonetheless, we believe this consistency can be facilitated if the MVM calculation uses a capital base consistent with pricing methods and practices and a cost of capital rate that is based on the rate used in pricing and which is calibrated as discussed below and in the full report.

This report examines two alternative approaches for defining the initial capital base in the MVM calculation, each of which uses a different risk exposure horizon, and the related implications for setting an appropriate cost of capital rate. These alternatives for defining the initial capital base are either currently used by insurers for product pricing or are beginning to be used for such purposes.

- An *ultimate risk exposure horizon* assumes that the relevant amount of capital for the MVM calculation is that which, at the start of each annual period, can support the risk in the ultimate liability cash flows. Some insurers' internal pricing methods and practices use this approach, since they often measure the capital base using required capital models developed by rating agencies, by regulators in North America under the NAIC's Risk-Based Capital standards and by the FSA under their Individual Capital Assessment Standards (ICAS).
- A *one-year risk exposure horizon* assumes that the relevant amount of capital for the MVM calculation is that which, at the start of each annual period, can support the risk of adverse changes in the best estimate liability in the coming year. Some insurers are beginning to use this risk exposure horizon in pricing practices because this is consistent with the solvency capital requirements for non-hedgeable risk under the Swiss Solvency Test or the proposed Solvency II framework.

Either of these approaches can be used to define the initial capital base for the MVM calculation. However, if the capital base is different using these two approaches then the cost of capital rate applied to each capital base would also have to be different so that the same MVM estimate is obtained.

- *<u>Pragmatic Approaches to Quantifying Risk Must Be Validated</u>*

  Given the complexity of the risk modeling required for many insurance products, practical and well-specified approximations are needed for actual implementation. For example, simplified methods for approximating the capital base needed in each future period of the MVM calculation and for approximating the impact of mortality and lapse risks are needed to make the calculations practical. Validation of these approximations is required to ensure that they produce risk measures that reflect the confidence levels and time horizons intended by the conceptual framework and that are consistent across risk categories and products.

- *<u>Cost of Capital Rate Must be Calibrated</u>*

  Some applications of the cost of capital methodology currently use a placeholder cost of capital rate of 6%. There have been few attempts, if any, to calibrate the resulting MVM estimates against observable benchmarks. We believe that the cost of capital rates used should be calibrated to ensure that they produce fair value liability estimates that are consistent with actual market prices for the assumption and transfer of insurance risks.

  The cost of capital rates that are obtained through this calibration exercise should also be assessed against objective criteria that reflect such factors as investor return expectations for insurance risk, the impact of corporate taxes and the impact of frictional costs of carrying capital. In the absence of such calibration against observed prices and objective criteria, it will be difficult for market participants to assess the reasonableness of any MVM calculations and their consistency with actual market pricing of insurance risks.

Implications

Our analysis of the cost of capital methodology indicates that in order to reliably calculate the fair value of insurance liabilities practitioners must be careful to consistently define the key inputs, validate pragmatic approximations used and calibrate the overall results to observed benchmarks. The specific challenges associated with these steps differ for the financial reporting and solvency applications addressed in the report.

Financial Reporting

For financial reporting applications, the fair value estimates should at all times reflect the current market environment with respect to the assessment of risk and the market price for the transfer of this risk. As a result, our findings with respect to the need to validate the approximations used in practice and the need to calibrate results against observed benchmarks should be carefully considered.

Implementation of the cost of capital methodology is likely to be facilitated if its parameters, the capital base and the cost of capital rates, are set consistently with market pricing practices. As noted above, this will require analysis of the methods and parameters used by insurers to price insurance and reinsurance products and capital market transactions involving insurance risk.

Solvency Applications

For solvency applications, the same implications apply. In addition though, solvency applications aim to quantify the potential *change* in fair value liabilities (in a distress scenario) over a specified solvency time horizon, such as one year in the proposed Solvency II and in the Swiss Solvency Test. This raises two unique challenges.

First, some practitioners advocate coordinating the time horizon for the solvency assessment (one-year) with the risk exposure horizon used to determine the capital base for the MVM calculation, resulting in the recommendation to use the one-year risk exposure horizon for MVM purposes. As previously noted, using this approach can lead to an appropriate MVM estimate. However, the cost of capital rate used may have to differ, perhaps significantly, from the cost of capital rates used by insurers whose internal pricing methods rely upon an ultimate risk exposure horizon.

Second, calibration of the capital base and the appropriate cost of capital rate are both more challenging in this context because of the need to reflect market conditions at the end of the solvency time horizon and after the occurrence of a “distress” event. As discussed in more detail in the report, the appropriate capital base for the MVM calculation and the cost of capital rate under these conditions may differ significantly from those implied by current market conditions. Failure to recognize these potential sources of change in the cost of risk may significantly affect estimates of required capital or available financial resources.

Conclusions

The cost of capital methodology presents a conceptually appealing and intuitive approach to determining market value margins for insurance risk. Nonetheless, we believe there are implementation challenges that suggest the need for extensive validation, calibration and consistency tests.

Given the significant role that this method is expected to play in emerging financial reporting standards and solvency guidelines affecting all insurers globally, we recommend that the parties responsible for establishing these financial reporting standards and solvency guidelines work swiftly, with the help of the industry, to perform these analyses, which might include the following:

- Comparison and analysis of the relationship of the ultimate and one-year risk exposure horizons for insurance risk to capital standards in common use, such as current rating agency requirements, U.S. NAIC RBC standards, Canadian MCCSR requirements and the FSA ICAS standards.

- Review of common pricing methods currently in use, focusing on the manner in which capital is allocated for pricing purposes and the rates of return or other pricing targets that are used in the process.
- Analysis of investor return expectations and the underlying capitalization levels inherent in such expectations.
- Review of the existing literature on the application of CAPM, the Fama-French 3-Factor Model and other means of establishing equity returns expected by investors in insurance companies with the goal of establishing a range of rates that can be considered consistent with investor expectations under alternative assumptions of the capital base.
- Development of estimates of factors, such as debt costs and the impact of corporate income taxes, that may impact the appropriate cost of capital rate.
- Review of capital market transactions, including securitizations of life and non-life insurance risks as well as non-insurance risks, to determine the implied cost of capital rate, as well as the potential change in this rate following a market shock or distress scenario.
- Recognizing the on-going debate over the appropriateness of using entry prices for calibrating fair value estimates, reconciliation of product prices to fair value liability estimates at the point of sale may shed light on the combination of risk measures and cost of capital rates needed to replicate market prices.

1. Introduction

1.1. Scope of Report

Two important insurance industry developments have brought considerable attention to the need for a consistent framework for measuring the fair value of insurance liabilities. First, new financial reporting standards aimed at producing *fair value* financial statements for insurers are being considered by the IASB, the FASB and other accounting policy organizations. Second, a variety of insurer solvency standards being considered in various countries have focused on the potential variability of fair value financial statements as the basis for establishing minimum regulatory capital requirements for insurers.

A key element of the calculation of fair values for insurance liabilities is the determination of a market value margin (MVM). While several methods of calculating MVM's for insurance liabilities have been explored, the *cost of capital method* has emerged as a strongly preferred methodology and has already been adopted in certain jurisdictions for the purposes of establishing capital requirements.

This report examines the cost of capital method in two specific contexts:

- Establishing the fair value of insurance liabilities for financial reporting purposes, and
- Establishing regulatory capital standards for insurers.

As the report will demonstrate, each of these applications can use the same underlying method for estimating the MVM for various insurance products. However, there are subtle yet important differences in the implementation of the cost of capital method under these two applications.

The report examines this method in detail and assesses the practical implementation issues associated with this approach to estimating risk margins. Specific examples of its application to life and non-life insurance products are used to highlight the issues most important to its appropriate implementation.

1.2. Outline of Report

The main elements of this report are summarized as follows:

- Section 2 provides an overview of the role of risk margins for ***financial reporting*** applications and an overview of the cost of capital method.
- Sections 3 through 5 explore the cost of capital method in the context of ***financial reporting*** and discuss a variety of challenges that the industry faces in ensuring that consistent and reliable results can be obtained.
- Section 6 discusses the role of risk margins in ***regulatory solvency*** frameworks and addresses additional aspects of the cost of capital method that are unique to solvency applications.
- Section 7 discusses the importance of the ability to properly calibrate the risk margins to ensure that the resulting risk margins adequately capture, along with the expected present value liability, the fair value of the liability.
- Various appendices contain additional technical details referenced in the main body of the report.

1.3. Additional Considerations

We acknowledge that, in addition to the implementation of this specific cost of capital method for the purposes of calculating market value margins, there are a variety of broader implementation issues and public policy considerations associated with both fair value financial reporting and solvency frameworks. Some examples include:

- Various issues regarding the determination of policy and contract cash flows to be included in liability estimates for financial reporting purposes;
- The appropriateness of standardized rather than company-specific risk margins for financial reporting;
- The procedures that might be required to ensure consistency between internal models and standard formulas;
- The appropriateness, under certain proposed solvency frameworks such as Solvency II and the Swiss Solvency Test, of a one-year solvency time horizon for setting total capital requirements on long-term insurance risks;
- Components of capital requirements under certain proposed solvency frameworks other than those that affect the calculation of risk margins, such as provisions for market, credit and operational risks.

Despite the interest in these issues among certain market participants, our focus in this report is solely on the practical implementation of the cost of capital method for estimating market value margins applicable to non-hedgeable insurance risks. We therefore do not specifically address these related issues.

2. Market Value Margins for Financial Reporting

2.1. Fair Value Financial Reporting Framework

While international accounting standards for insurers do not currently require fair value measurement for insurance liabilities, efforts are underway to move towards a *fair value* framework for financial reporting. Under a fair value framework, all assets and liabilities would be presented on a basis that reflects the price at which an asset can be sold to a third party or a liability could be settled with a third party.

In the specific instance of insurance liabilities, these are legally more difficult to literally transfer or settle than other financial liabilities. For these liabilities the fair value represents the cost of transferring the risk associated with the liability to a third party.

There are three components of the fair value of an insurance liability in this context:

- Best Estimate Liability

  A key component of the fair value for an insurance liability is an unbiased estimate of the expected cash flows on a present value basis. This value is referred to as the *best-estimate liability* (BEL).

  As we discuss in more detail below in Section 2.2, this definition assumes that the BEL does not incorporate a risk margin. However, in some instances it may be difficult to separate the true "best estimate" liability from risk margins associated with hedgeable market and credit risks.

  Market Value Margin

  The BEL, whether it includes a risk margin for hedgeable risks or not, will not fully reflect the price that would be paid to a third party to assume the risks associated with the liabilities. The price paid would also include a risk margin to compensate the third party for the non-hedgeable insurance risks as well, since these risks will generally be managed by holding capital. The risk margin is therefore needed to compensate the *capital providers* whose capital is used to absorb the risk assumed.

  Throughout this report this risk margin will be referred to as a market value margin (MVM). This terminology is used for consistency with the existing public discussions of these methods and is not meant to imply that there is, nor that there needs to be, an active, liquid market for trading the liability.

- Profit Margin

  A distinction can be made between the minimum risk margin *required* by the firm or its capital providers and the risk margin that is actually reflected in the premiums paid. This difference can be thought of as a *profit margin*.

  While the measurement of and accounting for the profit margin may be important in terms of the reported financial results at any point in time, for the purposes of this discussion there is no need to make this distinction. Throughout this discussion the premiums charged will be assumed to reflect only the minimum required risk margin, with the additional profit margin set to zero. This will allow us to leave issues such as the recognition of this profit margin in financial statements out of the present discussion.

The first two of these components are depicted graphically below relative to the total assets of the insurer:

Table 1: Fair Value Balance Sheet Components

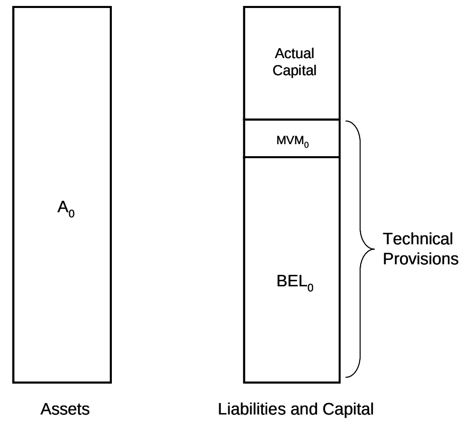

Here, the sum of the BEL and the MVM account for the total *technical provisions* in a fair value context. Since the objective is merely to report the financial condition, there is no explicit standard that the *capital* account must meet, so we depict the balance as the actual capital. Later, in Section 6, we will discuss how this framework can be adapted to also address capital standards for solvency regulation purposes.

2.2. Separation of the BEL and the MVM

In the discussion above, it was assumed that the BEL and the MVM could be determined separately. But in many instances the BEL for insurance liabilities may include a portion of the risk margin attributable to *hedgeable* risks and separation of those two quantities may not be feasible[1].

For this reason, it is often understood that the BEL will include the risk margin for all hedgeable risks related to equity prices, interest rates, exchange rates, etc. and that the MVM will reflect only the *additional* risk margin required for *non-hedgeable* risks related to mortality, lapse, policyholder behavior and other non-hedgeable insurance risks.

The proper accounting treatment of these explicit and implicit risk margins, including income recognition, balance sheet presentation and disclosure is subject to ongoing discussions since any differences between the treatment of the BEL and the MVM will result in inconsistent treatment of the risk margins for hedgeable and non-hedgeable risks.

While these accounting considerations may be important, for the present purposes we acknowledge that separation of the BEL and the risk margin for hedgeable risks is not feasible in many instances. This has implications for the ability to properly calibrate the MVM based on observable market transactions, which we will return to in Section 7.

Nonetheless, throughout this report we treat the MVM as representing the risk margin for non-hedgeable risks only.

2.3. Calculating the MVM for Financial Reporting Using the Cost of Capital Method

Risk margins play a critical role in the pricing of insurance and reinsurance policies, though practices for determining these risk margins continue to vary widely across different products, different markets and different insurers and reinsurers.

---

[1] For example, when insurance liabilities contain embedded interest rate or equity derivatives, risk neutral valuation methods that are commonly used to value the embedded derivative do not readily allow for separate quantification of the best estimate liability and the risk margin. Instead, these valuation methods produce only the aggregate *market value* of the embedded derivative.

One method for measuring risk margins, referred to here as the *cost of capital method*, has been strongly endorsed by various industry groups and regulatory bodies for estimating risk margins to be used in both financial reporting and solvency applications.

The cost of capital method is grounded in the basic premise that insurers hold capital to assure policyholders that their claims will be paid in the event that aggregate claim costs exceed the premiums paid. Because of the close connection between the amount of risk assumed and the amount of capital that must be held, risk margins can be thought of as representing the compensation required by the providers of that capital.

This leads to a conceptually appealing approach to establishing risk margins, with three key components:

- Capital Base – The first step is to assess the degree of risk in the underlying claim and benefit payments. From this, one can determine how much capital is needed to support the risk that these claim and benefit payments exceed their expected values[2]. This serves as the capital base for the MVM calculations.
- Time Horizon for Capital Commitment – The next step is to determine the period over which the capital must be committed and hence the capital base in *each period* until the final claim liability is paid. In the case when actual claim and benefit experience emerges as expected, capital intended to support the risk of deviations in these amounts from their expected value can gradually be released.
- Required Rate of Return on Capital Per Period – Finally, for each period during which the capital is held to support the risk associated with the unpaid claims and benefits, the providers of that capital will require an appropriate rate of return. Part of that return will be earned through the investment of the capital itself in marketable securities, so the *rate* of return used for this purpose need only reflect the spread over the risk-free return. The spread should reflect the amount of risk and may also include other "frictional costs" associated with holding capital, if any.

These three key inputs are used to determine the aggregate amount of compensation required, on a present value basis, by the capital providers.

This process is summarized graphically as follows:

Table 2: Demonstration of Cost of Capital Methodology

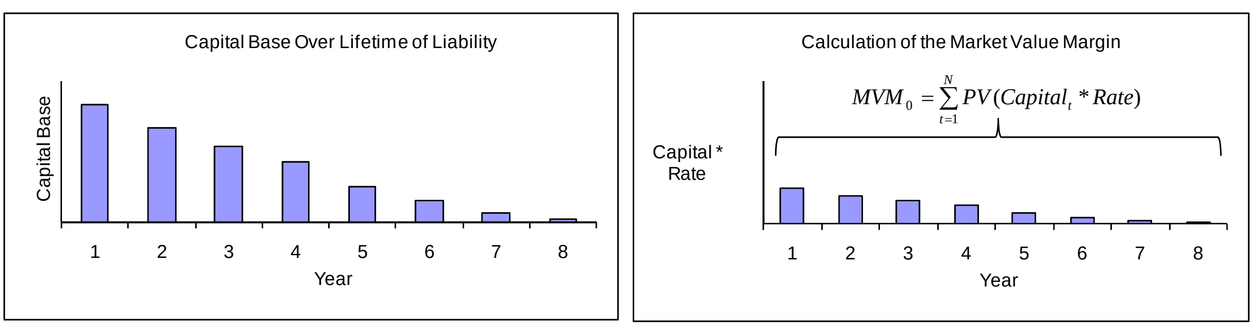

Numerical examples of this methodology will be used in the following sections of the report to highlight critical aspects of the calculations. An introductory numerical example is included as Appendix B for readers who are not familiar with the basic methodology.

[2] Because the cost of capital method is being discussed here only in the context of establishing risk margins to support non-hedgeable insurance risks, only these risks are taken into consideration when determining the degree to which payments can exceed their expected values.

3. Initial Capital Base

A critical element of the cost of capital method is the determination of the *initial* capital base used to reflect the risks faced by the party assuming the insurance risk. The complete calculation of the risk margin requires estimates of this capital based in all subsequent periods as well, which will be addressed fully in the next section. Here, we focus on the initial capital base only.

A number of factors can materially impact the measurement of the initial capital base, resulting in significant differences in reported MVM's. To ensure consistent application of this methodology, specific consideration must be given to each of the following:

- Link between the amount of capital and the annual cost of capital;
- Basis for initial capital base;
- Length of the risk exposure horizon used to measure risk at any point in time;
- Methods for measuring capital base to support non-hedgeable risks;
- Degree to which diversification benefits are reflected in the initial capital base.

3.1. Link Between the Amount of Capital and the Annual Cost of Capital

There are several alternative ways to define the initial capital base. These alternatives are, in their own right, important determinants of the MVM. But the *amount* of capital used in the cost of capital methodology is inextricably linked to its cost as a *rate per period* (as well as the time period over which the capital is held). All else being equal, for the same BEL a lower (or higher) amount of capital held should be used along with a higher (or lower) rate per period.

To demonstrate this, we abstract away from a specific insurance application and use the example of a put option on the S&P 500 instead. This "product" behaves much like an insurance policy but has the added feature that we can use the Black-Scholes option pricing formula to get a reliable estimate of a market consistent value for the option[3].

To keep this example consistent with insurance product examples, we assume a company is set up to write a single one-year put option on the S&P 500 index, raises capital to fund the potential losses on the transaction and invests both the capital and the put premium in a risk free bond earning 4%. We assume the current index value is 1,400, the strike price for the put is 1,400, the expected return on the index is 12% (no dividends), the risk free rate is 4% and the volatility is 15%. This produces a Black-Scholes put price of $57.51. The expected option payoff, our BEL, is $33.31 and the 99.5th percentile payoff is $339.31 given the assumptions here.

Now assume that the capital base needed to support the risk is given by the 99.5th percentile loss less the BEL, both on a present value basis, or $294.07. Applying the cost of capital method in this case involves applying the cost of capital rate for one period to this capital base. To match the known value of the option, the cost of capital rate would have to be 8.23%, as shown below:

[3] Notice that we do not reflect the impact of the issuer's own credit standing in this market value. This is consistent with our understanding of the current status of the on-going debate regarding whether or not fair value liabilities for insurers should or should not be adjusted to reflect the issuer's credit standing.

Table 3: Cost of Capital Method to Value Put Option – 99.5% VaR

Capital Base	294.07
Time Horizon	1.00
Cost of Capital Rate	8.23%
MVM	24.20
BEL	33.31
Price	57.51

Alternatively, if a different risk measure were used a different annual cost of capital rate would be required to obtain the same market consistent price. For instance, if the capital base were defined with respect to the 99.0% VaR instead of the 99.5% VaR, the rate would have to be 9.48%, as shown below:

Table 4: Cost of Capital Method to Value Put Option – 99.0% VaR

Capital Base	255.22
Time Horizon	1.00
Cost of Capital Rate	9.48%
MVM	24.20
BEL	33.31
Price	57.51

Although the cost of capital methodology discussed here is intended to be used to value non-hedgeable insurance risks rather than the risks depicted in this particular example, the calculations above effectively demonstrate the importance of ensuring consistency between the capital base and the rate applied to that capital. Since external market inferences for this rate (e.g. rates typically observed in the market) are likely to be used in practical implementation of the cost of capital methodology, using a capital base that differs from current market practice requires an adjustment to the rate used. This issue is addressed more thoroughly in Section 5.

3.2. Basis for Initial Capital Base

The alternatives for establishing the initial capital base include the following:

- Regulatory Required Capital – To some, the relevant capital base for the MVM calculation is the amount of capital required under a particular regulatory regime.

  One limitation of this approach for financial reporting purposes is that it could result in different values for the MVM for identical liabilities valued under different regulatory regimes, unless appropriate adjustments are made to the annual cost of capital rate. Although Solvency II may harmonize regulatory capital requirements in the EU, this challenge would remain for insurers operating outside of the EU. A consequence is that the sum of the best estimate liability and the MVM may only adequately capture the cost of transferring the risk associated with the liability to a third party operating within the same regulatory regime.

  Another limitation is the fact that regulatory models tend to be formulaic (often factor based) and have limited flexibility to account for risk differentiation across insurers or to reflect unique risk profiles. It is difficult to develop a comprehensive set of standard factors that can be applied to all firms, given the differences in product features, mix of business, mix of insureds, portfolio size, etc. For P&C insurers, the standard models would need to adequately capture the effects of different attachment point and policy limit profiles, different premium volumes and different policyholder profiles. For life insurers, the standard models would have to be capable of distinguishing between firms with unique product features or different risk management strategies embedded in crediting rate strategies, etc.

Sacrificing some degree of customization in favor of simplicity and consistency may be acceptable for solvency regulation purposes. Instances where the standard formulas do not adequately capture a specific insurer's unique risk profile can be addressed directly with the regulatory bodies. However, this may prove to be more problematic in a financial reporting setting, since third party users of the financial statements would expect an insurer's financial disclosures to accurately reflect the company's specific risks, products and risk modeling assumptions.

Furthermore, actual prices for transferring risk to third parties may not be driven entirely by the capital required for regulatory purposes. In many instances the capital actually held by insurers is higher than regulatory thresholds, with rating agency or internal capital models a more dominant driver of the capital held in practice.

- Rating Agency Target Capital – Many insurers operate under capital constraints dictated by rating agencies to a greater extent than those dictated by regulators. It would therefore be reasonable to assume that to the extent MVM's are driven by a desire to ensure an adequate return on capital, the relevant capital base would be the capital required by rating agencies for a given target rating level. In the event that the actual capital held by insurers is closely approximated by these rating agency capital targets, their use for the purposes of calculating the MVM using the cost of capital method may lead to a higher degree of consistency between the capital base and the annual cost of capital rate than would other definitions of the initial capital base.

  However, rating agency capital formulas or factors may not adequately account for differences in risk profiles across products or companies. As a result, it may be inappropriate for individual insurers to use these models as the basis for determining the fair value of their liabilities for financial reporting.

- Internal Capital Models – Many insurers have been developing internal risk models that are more customized to reflect company-specific risk profiles. These models can be useful in implementing a cost of capital approach to measuring the MVM. However, there are challenges to relying on internally developed models for the fair value financial reporting application addressed herein:

  - Implementation Requirements – Internal risk and capital models require extensive investments in computer software and require access to data and analytical expertise to properly perform model parameterization. While larger insurers likely have the resources to achieve this, for many smaller insurers these requirements may prove to be excessive.
  - Validation – Widespread use of customized, company-specific risk models introduces an equally difficult burden for auditors, regulators and others to validate model implementation, ensure completeness of the models and ensure that the probabilities assigned to specific scenarios are appropriate[4]. Experience to date with the FSA's ICAS and with Solvency II's quantitative impact studies suggests that significant unintended differences in implementation and results can occur from the wide range of plausible approaches.
  - Adjustments to the Annual Cost of Capital Rate – Often internal capital models may indicate a need for greater or less capital than what is actually held by insurers. When

[4] A case in point is the widespread use of *economic scenario generators* to derive stochastic forecasts of equity indices, yield curves and exchange rates. These scenarios are calibrated to be consistent with current market prices, however the results are not unique (i.e. many alternative scenario generators can all be consistent with current market prices) and the probabilities assigned to specific scenarios are not necessarily meaningful or easy to interpret.

this is the case, observed rates of return on actual capital held may need to be adjusted before they can be applied against the capital base indicated by the internal models.

Similarly, it may be the case that internal capital models lead some companies to carry more capital than the minimum regulatory thresholds or the target rating agency levels. This is because the models allow firms to focus on the goal of *maintaining* a target credit rating over a relatively long horizon, which may require more capital than merely attaining a current target credit rating. To the extent well-capitalized insurers choose to carry higher levels of capital to achieve this objective, their internal capital assessments may be inconsistent with the basis for the regulatory or rating agency models and adjustments to the annual cost of capital rate may be needed.

Regulatory requirements, rating agency targets and internal capital models may eventually converge, but presently they produce potentially significant differences in the capital base needed for the MVM calculation. These differences will lead to different estimates of the MVM unless adjustments are made to other key drivers of the risk margin, such as the annual cost of capital rate applied to these different capital bases.

3.3. Risk Exposure Horizon

One of the reasons why estimates of the capital base using regulatory, rating agency and internal models differ is that they often make different assumptions regarding which risks the capital is intended to absorb.

For instance, current NAIC regulatory capital models used for U.S. insurers and reinsurers assume that the initial capital must be sufficient to cover *ultimate* insurance liabilities, meaning that it must be sufficient to absorb deviations from expected results in all future periods. Existing rating agency capital models make a similar assumption[5]. In contrast, capital requirements under various existing and proposed regulatory capital frameworks specifically focus on capital requirements over a *one-year* solvency time horizon. And internal capital models, which at one time tended to mirror the rating agency approaches, are beginning to adopt a one-year horizon as well.

We believe the distinction between capital to absorb ultimate risk and capital to absorb one-year of risk at a time is critical to the determination of the MVM for insurance liabilities and requires careful coordination with the selection of the related annual cost of capital rate in order to be confident that market consistent results are obtained. To demonstrate this point we will focus in this section on the use of an internal capital model and avoid the nuances associated with specific regulatory or rating agency capital models.

Nonetheless, the points we raise are relevant to implementation of the cost of capital methodology for measuring the MVM even when specific rating agency or regulatory capital models are used. This is because those models, although they will commonly be more formulaic or rely on simplified risk factors, will inherently reflect one of these two alternative perspectives. As a result, to properly implement the methodology one would need to understand clearly which perspective is being used and make appropriate adjustments as described below.

3.3.1. Definition of Risk Exposure Horizon

The risk exposure horizon determines whether the initial capital base is intended to be sufficient, at the valuation date and at all subsequent dates, to absorb adverse deviations from expectations on an *ultimate*

[5] The percentile levels used in the various rating agency models are sometimes not clearly defined, though the S&P capital model for non-life insurers historically used capital factors derived from models of ultimate risk and measured using the equivalent of percentiles ranging from 96% to 99% depending on the line of business.

basis or, alternatively, whether the initial capital base is intended to provide capital sufficient to absorb adverse deviations from expectations just one year at a time, sequentially over the life of the business.[6]

In either case, the cost of capital method is intended to capture the costs associated with capital in *all* future periods over the complete time horizon over which claims and benefit cash flows are paid (i.e. the run-off period). As a result, in each case one can view the MVM as reflecting the cost of risks over the lifetime of the liabilities. However, in one instance (when the risk exposure horizon reflects risks on an ultimate basis) it is assumed that there is sufficient capital funded from inception so that future capital raising is not expected to be needed over the entire lifetime of the liabilities. In the case of sequential, one-year capital provisions it is assumed that capital can be continually and sequentially raised one year at a time as needed.

These alternative assumptions lead to different estimates of the capital base, although both capital definitions can be used to derive a market value margin using the cost of capital method by simply making appropriate adjustments to the cost of capital rate. We believe the approach using an "ultimate" risk exposure horizon is consistent with the method currently used by many U.S. insurers and reinsurers to price policies, such as those who rely on rating agency or NAIC regulatory capital models to define the capital base. We also believe that it is more consistent with current industry practice as it reflects the policyholders' expectations with regard to the financial strength of the insurer. The approach assuming that capital can be raised sequentially, one year at a time, is consistent with the method used in various proposals for setting capital requirements for regulatory purposes[7].

To clarify the important distinction between these alternative definitions of the risk exposure horizon, we use a simple numerical example to show how the alternative risk measures can be quantified and what the impact is on the full market value margin calculation. Later we will perform similar calculations for a more realistic insurance product.

3.3.2. Ultimate Risk Exposure Horizon

We begin with the risk exposure horizon measures on an ultimate basis and show how to quantify the risk in the case of a portfolio of 5-year term insurance policies. For this example we assume the following:

- 5-Year Term Insurance
- 1,000 policyholders, all the same age
- Fixed death benefit of $100,000 per insured life
- Mortality rates each period are independent and normally distributed with a mean of 2.0% and a standard deviation of 0.40%
- Fixed Premiums paid up-front

---

[6] It is important to recognize that the risk exposure horizon is often confused with other uses of the term *time horizon*. For instance, in the past there has been debate about whether the cost of capital method reflects only a single period of capital costs or whether it reflects capital costs in all future periods through the run-off of the liabilities. In all relevant applications of this method, and as we described in Section 2.3, it is clear that the method must include capital provisions in all future periods. Specific methods for doing so are addressed in detail in Section 4. In contrast, the risk exposure horizon defined here merely reflects whether, at each point in time, the capital base used within the MVM calculation is measured on an ultimate basis or measured sequentially, one year at a time.

[7] Two examples of this are the Swiss Solvency Test and Solvency II. In those frameworks, the cost of capital method is used in conjunction with a variety of additional considerations to establish minimum regulatory capital requirements. There are important differences in the implementation of the cost of capital method for these purposes as compared to the fair value financial reporting application discussed in this section. We will address these differences in detail in Section 6. We reference these applications here simply to highlight how the risk exposure horizon is defined in those specific applications.

- Risk-free yield curve assumed to be flat at 5% per annum (annually compounded) with no interest rate volatility over the 5-year period.

At the inception of this portfolio of risks, it is easy to determine the best estimate liability (BEL) immediately following payment of the premium as follows:

Table 5: Best Estimate Liability for 5-Year Term Life Example

Year	Qx	Lives - BOY	Deaths	Benefits Paid	Present Value Benefits Paid
1	2.00%	1,000	20.0	2,000,000	1,904,762
2	2.00%	980	19.6	1,960,000	1,777,778
3	2.00%	960	19.2	1,920,800	1,659,259
4	2.00%	941	18.8	1,882,384	1,548,642
5	2.00%	922	18.4	1,844,736	1,445,399
				BEL at t=0:	**8,335,840**

To estimate the potential for the ultimate liability to deviate from this BEL estimate, we can simulate the mortality rates in each period to derive the following distribution of results on a present value basis:

Table 6: Distribution of Present Value Liability – 5-Year Term Life

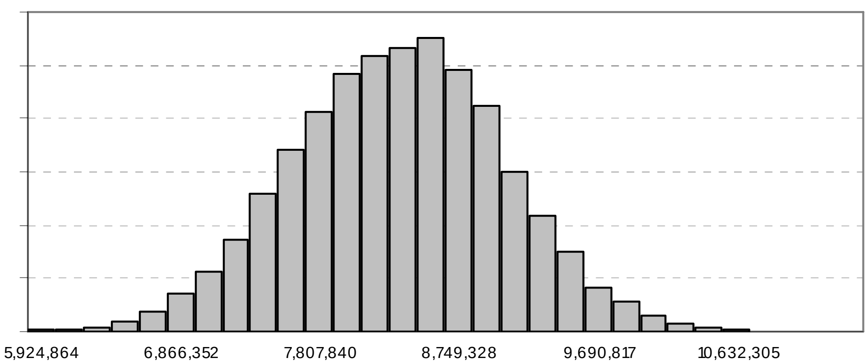

Count: 10,000 Min: 5,830,715 Max: 11,479,644 Mean: 8,339,162 Std Dev: 715,370

As this output shows, with 10,000 iterations we can approximate the actual mean present value benefit and see that the "worst case" in this simulation was a combination of mortality rates each period that led to present value benefits totaling approximately \$11.48 million. Although not shown explicitly, the 99.5th percentile is \$10.17 million, reflecting a potential present value deviation from the BEL of approximately \$1.83 million on an ultimate basis.

If a third party had a capital base of \$1.83 million and they assumed the risk associated with this liability, they would be 99.5% certain to have adequate resources to fully pay all claims. This risk measure reflects the potential variability in mortality rates in all future periods and therefore reflects what we refer to as an *ultimate risk exposure horizon*.

3.3.3. One-Year Risk Exposure Horizon

As an alternative to the previous approach to measuring risk, we show a comparable measure using a *one-year risk exposure horizon*.

In the most literal interpretation of incorporating risk over a one-year horizon, we would estimate the potential impact of an adverse mortality rate for the first period only, with all subsequent periods' mortality rates reflecting their *expected* values. Thus, over a one-year risk exposure horizon the mortality rates could be as high as the 99.5th percentile of the mortality rate distribution, or 3.03%. That would make the first year benefit payment significantly higher than its expected value but would result in approximately the same *expected* benefit payments in subsequent years (but for the impact that the higher than expected first period mortality rate has on the size of the insured population as of the end of the first period).

This would result in the following calculation of the best estimate liability conditional on a "worst case" mortality rate in the first period:

Table 7: Single Period Risk – 5-Year Term Life

Year	Qx	Lives - BOY	Deaths	Benefits Paid	Present Value Benefits Paid
1	3.03%	1,000	30.3	3,030,332	2,886,030
2	2.00%	970	19.4	1,939,393	1,759,087
3	2.00%	950	19.0	1,900,605	1,641,814
4	2.00%	931	18.6	1,862,593	1,532,360
5	2.00%	913	18.3	1,825,342	1,430,203
				Conditional BEL at t=0:	**9,249,495**

This calculation shows that if the first period mortality rate is significantly higher than its expected value, the present value liability (at inception) could be approximately $913,655 higher than the BEL.

Notice that this quantity is approximately half the value of the risk measured using the ultimate risk exposure horizon ($1.83 million). The larger number reflects the potential for future periods' mortality rates to also deviate from their expected value. The smaller amount considers only the risk that could materialize over a one year period and ignores the risk in the future periods[8].

3.3.4. Two Alternative Perspectives on Risk Exposure Horizon

The above numerical example shows that the ultimate and one-year risk exposure horizons may produce materially different risk measures. If each of these amounts were used to establish the initial capital base, the cost of capital method would also produce materially different measurements of the market value margin unless adjustments were made to the cost of capital rate applies.

The reason these differences arise is that the two approaches adopt different perspectives regarding when the risks need to be funded or collateralized:

- Ultimate Risk Must Be Fully Funded at Inception – This perspective assumes that the party assuming risks that can occur over long horizons must be adequately capitalized *from inception* to absorb all of the risks that could materialize over the full run-off period of the liabilities. It is equivalent to assuming that the risks need to be fully funded at inception even though some of the risks cannot materialize until later periods. This eliminates the need to raise additional funds unless the perception of future risks changes adversely.

- Risk Can Be Funded Sequentially, One Period at Time – The alternative perspective assumes that capital can be raised sequentially to fund a series of one-year risks. As a result, risks that cannot occur until later periods do not result in the need for capital at inception.

In what follows, we emphasize the difference in these two perspectives using a numerical example of the calculation of the MVM needed for the estimation of the fair value of the liability at inception. We also

[8] The two alternative risk measures use different risk exposure horizons. This may justify using different percentiles for each and the resulting differences may not be as significant as shown here. The example is being used simply to show the differences in the risks captured rather than to argue that any particular percentile should be used. The impact of using different percentiles for the ultimate risk exposure horizon calculation is shown in Section 3.3.4.

identify the critical differences in the application of the cost of capital method that each perspective requires and draw some comparisons to existing frameworks.

Determining Capital Base Using Ultimate Risk Exposure Horizon

Earlier we showed the initial measurement of the capital base using an ultimate risk exposure horizon for the 5-year term insurance example. Although computationally challenging in practice, it is possible to measure the risk at inception, as in the calculations above, as well as at any future point in time.

The table below recreates the original analysis at inception and then repeats that analysis at the end of each of the next four years assuming that actual experience to date has been as expected at inception[9]. For each period, a risk measure reflecting the difference between the 99.5th percentile present value of ultimate benefits and the expected present value of ultimate benefits is calculated.

Table 8: Capital Base Over Time Using Ultimate Risk Exposure Horizon

Year	Best Estimate Liability	Percentile	Liability at 99.5th Percentile	99.5th Percentile Capital Base
1	8,338,097	99.5%	10,165,700	1,827,604
2	6,755,724	99.5%	8,421,047	1,665,323
3	5,134,530	99.5%	6,604,840	1,470,311
4	3,467,711	99.5%	4,722,393	1,254,682
5	1,757,760	99.5%	2,665,175	907,415

This example shows four separate effects through time. One is that the exposure to risk declines in this case because the number of lives remaining in the portfolio declines. The second is that the potential variability from the expected benefit payments declines over time because there are fewer periods over which the mortality rates can deviate from their expected values. The third is that the impact of the decreased variability each period is dampened over time by the elimination of diversification benefits across periods. And the fourth is that the capital base reflects less discounting for the time value of money (the amounts shown above reflect present values at the start of each period rather than at inception).

In these calculations, at each point in time the remaining risk reflects the combined impact of potential deviations of actual from expected mortality rates in all future periods – an *ultimate* risk exposure horizon.

In addition, as was noted above, each of these calculations uses the 99.5th percentile as the basis for measuring the risk of adverse deviations from the best estimate liability, even though the calculations involve risk exposure horizons longer than just one year. Some practitioners argue that when longer term risks are considered the percentile should be adjusted downward. The basis for this is that the 99.5th percentile is often thought to represent the complement of the one-year default rate for corporate bonds. It is therefore argued and that when the risk is measured over longer risk exposure horizons the *n*-year default rate for bonds should be used instead of the one-year default rate. This higher default rate results in a *lower* percentile being used. While we will not address the merits of this particular aspect of the methodology, it is important to recognize that making such modifications to the percentiles would alter the calculations but would not alter the important distinction being made here between an ultimate risk exposure horizon and a one-year risk exposure horizon.

[9] Each of these *future* calculations is necessarily conditional on experience to date. Here, we assume that experience to date is always equal to the expected experience and that no assumptions are changed other than the remaining time horizon.

For example, if the calculations described above were redone using the following "default rates" for different time horizons[10,11]:

Table 9: Target Default Rates by Time Horizon

Years	Fitch BBB Default Probabilities	Target Percentile
1	0.5%	99.5%
2	0.7%	99.3%
3	1.2%	98.8%
4	1.7%	98.3%
5	2.3%	97.7%

then the following alternative results would be obtained:

Table 10: Capital Base Over Time Using Term-Specific Percentiles

Year	Best Estimate Liability	Target Percentile	Liability at Target Percentile	Term-Specific Capital Base
1	8,338,097	97.7%	9,755,330	1,417,233
2	6,755,724	98.3%	8,127,189	1,371,465
3	5,134,530	98.8%	6,430,006	1,295,476
4	3,467,711	99.3%	4,663,333	1,195,622
5	1,757,760	99.5%	2,665,175	907,415

Using the target percentiles that vary based on the term of the liability results in small differences in the capital base. These differences would have been even smaller had the weighted average term of the outstanding liabilities been used.

Determining Capital Base Using One-Year Risk Exposure Horizon

The alternative approach, using a one-year risk exposure horizon, is to estimate a series of sequentially determined one-year risk measures. The example starts by estimating how much the first year's benefit payments, and subsequent best estimates of the outstanding liabilities, could deviate from their expected values at the 99.5th percentile. This calculation was shown previously in Table 7. When this amount is compared to the current BEL, the first period's capital base is obtained.

Then, conditional on the mortality experience to date being as expected (as opposed to the 99.5th percentile tail result), we estimate for the second period how much the benefit payments could deviate from their expected value, again at the 99.5th percentile. This is continued each period to produce five separate sequential, one-year risk measures. For consistency with the previous numerical example, we measure each period's risk as of the start of each period rather than discount the amounts back to inception. We also show the sum of the amount paid in the relevant period and the expected outstanding payments to be made in subsequent periods (all discounted to the start of the period). The results are as shown below:

---

[10] The source for these particular rates is Fitch's, *Exposure Draft: Prism – Insurance Rating Calibration Measures*, except for the one-year rate which was set to 99.5% (instead of 99.7%) for consistency with the other calculations shown in this section.

[11] The particular adjustments to the percentiles for longer horizons are for illustrative purposes only and other adjustments may be needed. For instance, default rates based on corporate bond data may provide an inaccurate estimate of the probability of economic insolvency for an insurer (see Grondin, *Economic Capital Models and Implied Ratings*). In addition, the modified percentiles used here reflect the *term* of the liability rather than the average time until payment, which is significantly shorter than the term in this particular example. Most practical applications would more likely use the percentiles associated with the average time to payment, resulting in smaller differences between these results and those using the 99.5th percentile. In subsequent examples in this paper we will follow this approach and select the percentiles based on the average time until payment rather than the term.

Table 11: Sequential Calculation of One Year Risk Measures

Year	Expected PV Benefits Paid + Unpaid	Tail-Event PV Benefits Paid + Unpaid	Risk Measure
1	8,335,840	9,249,495	913,655
2	6,752,632	7,662,906	910,274
3	5,130,264	6,037,969	907,705
4	3,465,977	4,371,947	905,970
5	1,756,892	2,661,982	905,091

The results from Table 8, Table 10 and Table 11 are compared below to highlight the difference in these calculations:

Table 12: Comparison of Alternative Risk Measures

	Ultimate Risk Exposure Horizon		One Year
Year	99.5th Percentile	Term-Specific Percentile	Risk Exposure Horizon
1	1,827,604	1,417,233	913,655
2	1,665,323	1,371,465	910,274
3	1,470,311	1,295,476	907,705
4	1,254,682	1,195,622	905,970
5	907,415	907,415	905,091

Notice in each period other than the final period[12] the one-year risk measure understates the (remaining) ultimate risk because it does not consider the risks in any periods other than the immediate one-year period.

Applying the Cost of Capital Method

These alternative perspectives lead to materially different estimates of the capital base needed each period to support the risk. In one case, the assumption is that all of the risk must be fully funded at inception. In the other case, the assumption is that the risks can be funded sequentially, one year at a time.

We will address specific issues with various methods to efficiently estimate the capital base at the start of each subsequent annual period in Section 4 and specific issues with the factors that determine the appropriate annual rate of return to apply to these amounts in Section 5. Nonetheless, we can demonstrate the complete calculation of the MVM using the three alternative series of capital amounts using a 6% annual *cost of capital* rate and an assumed 5% risk free discount rate. These three calculations are shown as follows:

Table 13: Comparison of MVM Calculations (6% Cost of Capital Rate)

	Ultimate Risk Exposure Horizon @ 99.5th Percentile			Ultimate Risk Exposure Horizon @ Target Percentile			One Year Risk Exposure Horizon		
Year	Capital Base	Annual Cost	PV Annual Cost	Capital Base	Annual Cost	PV Annual Cost	Capital Base	Annual Cost	PV Annual Cost
1	1,827,604	109,656	104,434	1,417,233	85,034	80,985	913,655	54,819	52,209
2	1,665,323	99,919	90,630	1,371,465	82,288	74,638	910,274	54,616	49,539
3	1,470,311	88,219	76,207	1,295,476	77,729	67,145	907,705	54,462	47,047
4	1,254,682	75,281	61,934	1,195,622	71,737	59,018	905,970	54,358	44,721
5	907,415	54,445	42,659	907,415	54,445	42,659	905,091	54,305	42,550
		MVM:	**375,864**		MVM:	**324,445**		MVM:	**236,065**

Notice that while either perspective could be used *mechanically* to derive the MVM, the cost of capital *rate* per period would have to be adjusted in order to produce the same MVM. Without any adjustments to the cost of capital rate, for instance if 6% were used in each case, the one-year risk exposure horizon produces an MVM of \$236,065 compared to an MVM of \$375,864 in the case of the ultimate risk exposure horizon using the 99.5th percentile. These differences can be eliminated by adjusting the cost of capital rate. For instance, the one-year risk exposure horizon could be used with an annual cost of capital

[12] The modest difference in the final period amounts reflect the fact that the amounts in Table 8 are based on the results of a Monte Carlo simulation whereas the results from Table 11 are based on analytical calculations.

rate of 9.55% to produce the same MVM as the calculations using the ultimate risk exposure horizon at the 99.5th percentile and a 6% annual cost of capital rate.

Comparison to Existing Pricing Practices

Since the market will provide only one *fair value* liability price point, the capital base and the cost of capital rate that is applied to this base must be consistently derived. This process will be facilitated if both assumptions can be reconciled to existing pricing practices among insurers and reinsurers.

While we recognize that practices vary considerably among insurers and reinsurers, we note that many insurers and reinsurers use methods similar to the cost of capital method to develop risk margins for actual transactions, including methods that seek to achieve a target internal rate of return (IRR) on capital flows. Often existing regulatory, rating agency and internal capital models, which are generally more consistent with a view that the initial capital must account for the ultimate risk, are used for these calculations. For instance, under the NAIC's C-3 Phase II methodology the initial capital reflects the average of the worst shortfalls that exist in *any* future period, which is impacted not just by risks that can manifest in one year but also by risks that may materialize in future periods as well. Similarly, existing rating agency capital models reflect risk on an ultimate basis, as opposed to a one-year basis. Many existing internal capital models use a similar approach. As a result, the use of an ultimate risk measure for the MVM calculation may facilitate adjustment of the cost of capital rate used for pricing purposes so that the cost of capital method will produce fair value liability estimates that are consistent with market prices.

At that same time, emerging solvency capital requirements under the Swiss Solvency Test and Solvency II have led some insurers to adapt their own internal capital models to measure risks over a one-year horizon. Because of the convenience of using the same risk measures for solvency assessment purposes and for establishing the MVM, some insurers may adopt a one-year risk exposure horizon for pricing purposes as well. If this occurs with little or no change in prices, then the adjustments to the cost of capital rate demonstrated earlier would be required in order to obtain the same MVM under either approach.

In addition, when a one-year risk exposure horizon is used to calculate the MVM there is an implicit assumption that capital can always be raised sequentially, as needed, at a known cost. Consideration will have to be given to whether this is a reasonable assumption or if additional adjustments would have to be made so that the MVM calculation will produce fair value liability estimates that are consistent with market prices.

3.4. Measuring Capital Base for Non-Hedgeable Risks

In the simplified numerical example used in the previous section, the capital base using either the ultimate risk exposure horizon or the one-year risk exposure horizon was derived using simple analytical formulas or relatively trivial sequential simulations. In more realistic cases that would normally be of interest, insurance product features would likely make it far more difficult to reliably and efficiently estimate the capital base using either of these approaches. Instead, more practical *stress testing* is likely to be performed using various "shocks" to key assumptions.

We will not address the details of the various methods that may be needed for all products. We note, however, that these simplifications, which are necessary for practical implementation of the cost of capital method, will require extensive validation to ensure that the magnitude of the shocks and the nature of how their impact is measured results in consistent estimation of the capital base at the desired level of confidence (e.g. 99.5th percentile).

In Appendix C, we discuss one particular case where this may be more difficult than is sometimes recognized using a simplified version of an SPDA product that has exposure to both hedgeable and non-hedgeable risks. As mentioned briefly in Section 2.2, typically the estimate of the "best estimate liability" will already reflect an adjustment to reflect a risk margin for the hedgeable risks through the use of risk-neutral interest rate scenarios. As a result, the capital base for the MVM calculation should reflect only the additional capital needed to absorb the non-hedgeable risk. Care must be taken to avoid over- or underestimating the capital base in this case.

As shown in Appendix C, if the analysis of the impact of the non-hedgeable risks is made using a static set of future interest rate scenarios (such as the forward rates) as opposed to a particularly adverse set of interest rate scenarios, the impact of the non-hedgeable risks can be severely understated. In addition, if specific "shock" scenarios are used to measure the impact of specific non-hedgeable risks, calibrating these shocks to reflect the same level of confidence as a full stochastic simulation of the underlying non-hedgeable risks may be difficult to achieve.

3.5. Diversification Adjustments

Diversification adjustments play a critical role in calculating the initial capital base for MVM purposes, including adjustments to reflect diversification across different sources of risk, diversification across policies within a portfolio for a given product or line of business and diversification across products. Because the diversification effects will likely vary considerably for different products and for different insurers, consistency of results is difficult to obtain.

Of particular importance is the definition of the *reference entity* that should be assumed to represent the third-party who would assume the insurer's entire portfolio of risks. In theory, the risk margin required by this party would reflect its <u>existing</u> portfolio of risks and not just the diversified portfolio of risks being (hypothetically) transferred. That is to say, the party assuming the risk would take into consideration the further diversification of risk that occurs when the new risk portfolio is added to their existing risk portfolio.

In practice, it is unlikely that such diversification with the (hypothetical) reference entity portfolio can be quantified. Instead, only diversification that occurs *within* a given insurer's risk portfolio can realistically be reflected.

4. Time Horizon for Capital Commitment

Determining the initial capital base for the MVM calculation is a critical aspect of the cost of capital method. Once this is done, it is equally important to assess the time horizon over which the capital will be held to support the risk, or equivalently, what the capital base will be in each future period over the life of the liability.

The two key issues to address are how to approximate the capital base over time and how to reflect the impact of diversification.

4.1. Approximating the Capital Base Over Time

To apply the cost of capital method, assumptions have to be made with regard to how the initial capital base changes over time. In Section 3.3 we showed the pattern of capital over the life of the liability under two alternative perspectives with regard to when the risks need to be funded – at inception for the ultimate risk exposure horizon or sequentially for the one-year risk exposure horizon.

In each of those calculations the simplicity of the example enabled us to compute the capital base under either measure in each future period. In practice this is often more difficult than in the simplified example and approximations are needed.

4.1.1. P&C Product Applications

For P&C products it is somewhat natural to assume that the uncertainty is strictly *reduced* over time since the exposure to risk does not increase over time as the result of continuing premium payments or increases in account values as in many life insurance and annuity products. As a result, if the initial capital is based on an ultimate risk exposure horizon it is logical to assume that this initial capital can be gradually "released" each period.

<u>Capital Base Over Time Using Ultimate Risk Exposure Horizon</u>

A special case of the calculation of the capital base in all future periods is shown in Appendix B. There it is assumed that none of the uncertainty is resolved until the final payment of the claim and so the risk and associated capital base remains relatively constant each period. The only changes that occur each period are due to changes in the amount of discounting.

In a more general case, there are several ways to model the release of the initial capital base over time:

- Proportional to Loss Reserve Release – It is common to assume that the risk is resolved in proportion to the release of loss reserves (i.e. the payment of claims) and this gives rise to a natural amortization schedule for the initial capital base for most lines of business.

- Proportional to Establishment of Reserves – In some cases, practitioners argue that more relevant than the timing of the claim *payments* is the timing of when the insurer is able to accurately assess the unpaid claims, which may occur more rapidly than actually paying the claims. The suggestion is that once information is known about the magnitude of the claims, whether they are paid or not, capital not needed to pay the claims can be released. This gives rise to an accelerated capital release schedule relative to the baseline payment pattern.

- Implied by Rating Agency Capital Models – Another approach is to infer a *typical* release pattern through sequential application of a factor-based rating agency capital formula, such as the one used by S&P. In the S&P formula, capital is required for both underwriting risk and reserve risk (in addition to other risks that are not relevant to this particular discussion). When specific factors are applied to the written premium and the reserves at any point in time, they imply a particular pattern of release of the initial S&P capital base.

The assumptions used with respect to how the capital can be released over time can have a significant impact on the MVM calculation. Although the mechanics are trivial, little empirical analysis exists regarding the timing of the release of *uncertainty* in P&C liabilities. Company and industry data is less useful for this purpose than it may appear, in part because of common practices associated with bulk reserving and the absence of commonly used methods to report measures of uncertainty over time. Historical estimates of expected values are insufficient to accurately quantify the timing of the release of capital, except using *ad hoc* adjustments.

Capital Base Over Time Using One-Year Risk Exposure Horizon

If the one-year risk exposure horizon perspective is taken, it isn't clear that any of the methods described above for determining the capital base in future periods are appropriate. Instead, specific assumptions regarding the magnitude of the one-year risk each period would have to be made.

This point is demonstrated more fully in Appendix D, which shows an example using a General Liability product and assumptions derived from industry data. Using the same approach as outlined in Section 3.3, which there reflected a simplified mortality product, the MVM using an ultimate risk exposure horizon is calculated as being equal to either 12.3% of the best estimate liability if the 99.5th percentile is used or 9.4% of the best estimate liability if the percentile is adjusted to reflect the average time to payment of 4.2 years. However, if the risk exposure horizon is adjusted to reflect solely the risk in the first year and no adjustments are made to the annual cost of capital rate, the MVM declines to only 5% of the best estimate liability.

Part of the reason for this large discrepancy is the fact that it becomes far more difficult to accurately depict the series of one-year capital measures using only the one-year risk exposure horizon amounts and the pattern of changes in the best estimate liability. This is true for many lines of insurance, particularly high-layer excess of loss reinsurance on a long-tailed liability risk such as medical malpractice or product liability. In the first year there is often very little learned about the frequency or severity of claims against these policies. As a result, it is common for the estimates at the end of the period to differ insignificantly from the estimates at inception. For instance, as shown in Appendix D, the one year volatility of the loss ratio using industry General Liability data is only 6.8% while the volatility of the ultimate loss ratios for the same 9-year historical period is 15.1%. Whereas the one-year risk exposure horizon picks up some of the potential adverse experience, at some point in time the sequence of capital based on a series of one-year risk exposure horizons will result in the need for additional capital commitments. As a result, methods that simply "release" the initial capital base proportionately cannot accurately capture the sequence of one-year risk measures.

A similar result for a portfolio of loss reserves is shown in Appendix E. There we use a stochastic loss reserving model to produce estimates of the variability of ultimate loss estimates in each year over the run-off of the portfolio. Using the model results for a portfolio of commercial auto liability reserves and a portfolio of general liability reserves, we show that an approximation using the initial estimate of the capital base for a one-year risk exposure horizon and the ratio of outstanding reserves at each point understates the MVM significantly. The approximation results in an estimated MVM for the commercial auto and general liability portfolios that are only 62% and 36%, respectively, of the MVM that would be produced by actually modeling the sequence of one-year risk measures directly.

4.1.2. Life and Annuity Product Applications

For some life insurance and annuity products, the growth in account values from new premium deposits or from investment returns could result in larger risk measures and associated capital amounts over time. As a result, it is common to assume that simple risk drivers such as statutory reserves, account values or some other risk exposure base can be used to determine the relative capital levels in each period.

In many cases though, this may be inadequate. Consider the 5-year term life example discussed in Section 3.3. There, we noted that the use of an ultimate risk exposure horizon results in a risk measure that declines as time passes. In the first year of a 5-year term insurance policy (for a large cohort of policyholders) we are concerned with adverse mortality experience in each of the next five years. In the last year though, we are only impacted by the risk of adverse mortality experience in this final year. Therefore, the risk measure is not simply proportional to an exposure base such as the number of insured lives or the BEL (reserve), as shown in the table below:

Table 14: Comparison of Capital Base to Risk Driver Over Time

Year	Best Estimate Liability	Capital Base	Capital to BEL Ratio	BOY Insured Lives	Capital per Insured Lives
1	8,338,097	1,827,604	22%	1,000	1,828
2	6,755,724	1,665,323	25%	980	1,699
3	5,134,530	1,470,311	29%	960	1,531
4	3,467,711	1,254,682	36%	941	1,333
5	1,757,760	907,415	52%	922	984

As this example shows, the capital grows as a percentage of the BEL and declines relative to the number of insured lives. The same issues would apply in a more complex manner in cases where the exposure grows with time.

4.2. Impact of Diversification

Another challenge to address is how to reflect changes in diversification benefits over time.

As discussed in Section 3.4, the initial capital base is impacted by adjustments to account for diversification across risks, across policyholders and across products. Over time, some of these diversification effects are likely to disappear as the mix of outstanding risks changes (e.g. when short-term liabilities mature and only certain longer-term liabilities remain). This could result in a significant increase in the capital base needed to support any one product relative to its initial capital base.

For instance, consider an insurer initially writing short-tailed and long-tailed risks, with claim payments in 2 years and 10 years respectively. Each of these two separate portfolios is assumed to require $100 million of capital on their own, but they require only $150 million in total due to the diversification across these products. Further assume that we have concluded that the diversified capital base for each product is $75 million and that it is “released” over the payment horizons of 2- and 10-years, respectively. Notice that after 2 years, the short-tailed line will no longer exist, having paid all of its benefits, and the portfolio will consist solely of the 10-year risk. At that moment, the diversification effect can no longer be assumed to apply. The capital base in each subsequent period therefore cannot simply reflect a linear amortization of the initial capital base. Instead, it must reflect an increase at the end of year 2 when the diversification effect no longer applies.

5. Annual Cost of Capital

The final input to the cost of capital methodology is the annual rate of return applied to the capital base in each period.

Current applications of the cost of capital method have not yet fully addressed how this rate is to be determined. Solvency II's quantitative impact studies and the Swiss Solvency Test, which both use versions of this methodology for establishing regulatory capital requirements, currently use a placeholder rate of 6% per year. In other contexts, such as European Embedded Value reporting, some insurers have used lower rates.

It is possible to implement the cost of capital methodology in a financial reporting framework by simply setting this rate to an arbitrary fixed value. However, we do not believe that this can provide any assurance that the resulting MVM is adequate to compensate a third-party for the risk unless the rate is well defined, the rate is appropriate for all products and all firms, and the rate is carefully calibrated to actual market prices for insurance risk. As a result, it is important to either benchmark the rate used against objective standards or calibrate the final cost of capital MVM against observable prices in the marketplace.

The following considerations are explored in this section:

- Basis for Establishing the Cost of Capital Rate
- Impact of Corporate Income Taxes
- Calibration to Market Prices

5.1. Basis for Establishing the Cost of Capital Rate

5.1.1. Total Return vs. Spread

The annual rate used in the calculation of the MVM under the cost of capital method has to be sufficient to ensure that there is an adequate pool of funds to compensate the capital providers for the risk that they assume. Because the capital itself is assumed to be held in marketable securities, the MVM does not have to provide for the *total* return and instead merely has to provide for the *spread* over and above the risk free rate of return. For convenience in the discussion that follows, we will refer to this quantity as the *cost of capital rate*.

5.1.2. Return Requirements of Capital Providers

The "return on capital" framework that is used in the cost of capital method makes it natural to assume that existing models of the returns demanded by capital providers in other markets could be used to derive objective measures of the rate that should be used in the MVM calculation.

For instance, shareholder return models such as the Capital Asset Pricing Model (CAPM), including variations and extensions thereof, can be used to quantify the risk premium (the rate of return per period above the risk free rate) that shareholders of insurance companies demand to assume broadly diversified insurance risk. In some instances, this rate may be appropriate for use in the calculation of the MVM, though several adjustments are likely to be needed.

<u>Applicability of Shareholder Based Models</u>

Despite the intuitive appeal of relying on shareholder-based models to estimate expected costs per unit of (risk) capital each period, commonly used shareholder return models are perhaps less useful than they may first appear. One reason is that they use a different definition of "risk" than used here; another reason is returns for the shareholders of a going concern are not necessarily the same as those of a run-off insurer.

- Risk Definition – Models such as the CAPM or the Fama-French 3-Factor Model are commonly used in a variety of investment and corporate finance applications. Part of their appeal is that they attempt to specifically quantify the excess return over the risk free rate that shareholders expect to earn given the level of risk assumed. However, they reflect very specific definitions of "risk" that are sometimes incompatible with the definition of risk used in the risk and capital models discussed previously. Specifically, they focus entirely on *systematic* or non-diversifiable risks and typically argue that these are the only risks for which investors are compensated.

  Despite substantial theoretical support for these models, their results are often inconsistent with observed pricing for many risks. This is particularly true for insurance risks where the systematic risk is believed to be quite low. For instance, property-catastrophe bonds are often thought to contain minimal systematic market risk yet their prices reflect significant risk margins.

  Appendix F discusses a variety of practical issues associated with applying these models to determine the expected shareholder returns for insurance companies. There, we show that there is significant variation in estimates of the critical risk measure in the CAPM (the insurance *beta* relative to the broader market) due not only to differences in risk profiles across insurers and industry sectors, but also due to measurement errors and related statistical issues. In addition, the range of estimates for the market risk premium is quite large, depending upon the time period used and how the historical data is analyzed. As a result, CAPM estimates of the risk premium (over the risk free rate) for P&C insurers can range from 2.01% to 8.65%.

  The Fama-French 3-Factor Model includes two additional sources of systematic risk premiums beyond the market risk premium reflected in CAPM. As a result, risk premiums for P&C insurers can exceed 14% when industry data is used to parameterize the Fama-French model.

  Because these risk premiums are applied to the capital base *per annum*, for many insurance lines of business or products the differences in the resulting MVM estimates can be significant. In the example shown in the appendix, the MVM ranges from $2.57 million to $17.30 million as a result of differences in the selected risk premium.

- Returns for Run-Off/Closed Block – In the fair value financial reporting applications (as well as the solvency applications discussed in Section 6) the MVM is limited to existing liabilities. This may imply an assumption that the portfolio would be immediately placed into run-off. As a result, total shareholder return models may need to be adjusted to account for the fact that a key source of "return" that exists for going concerns does not exist and this may alter the returns demanded by the capital providers[13].

<u>Role of Corporate Debt and the Weighted Average Cost of Capital</u>

The CAPM and Fama-French models discussed above are used to measure the return expectations of shareholders. For insurers who are able to fund part of their capital base through the issuance of debt, we often observe return on capital targets being set at their weighted average cost of capital (WACC) rather than simply the cost of equity capital. To the extent this is done, the average rate will be lower than the pure equity return rate.

An important issue to address in this regard is the extent to which the MVM under the cost of capital method should take the existence of debt into account. If the capital base used in the calculation can be funded with a mixture of debt and equity capital, then the WACC should be used. However, we question whether this is an appropriate assumption. Recall that the capital base is typically defined in terms of the amount of capital that is substantially at risk (e.g. it reflects a diversified measure only up to a specified

---

[13] See Exley and Smith for a more complete discussion of this issue.

percentile). It seems inconsistent to consider this capital substantially at risk and also assume that it can be funded by debt investors at costs substantially below the equity costs.

Frictional Costs

Some researchers assert that there are additional "costs" for an insurer carrying capital, beyond those required to compensate investors for the risk they are assuming. These additional costs reflect a variety of indirect costs that are difficult or impossible to quantify, such as frictional costs related to managers' incentives or information asymmetries. Despite the theoretical arguments supporting this notion, we do not believe that these can be reliably quantified to the extent needed for the purposes discussed here.

A somewhat more direct frictional cost of holding capital results from the double taxation of investment income, first when earned by the insurer and then when paid to the shareholder in the form of dividends. Assuming that this cost reflects the effective tax rate on the risk free income on the capital only (ignoring taxes on other sources of investment income), this cost could add an additional 1% or so to the annual spread[14].

For the applications addressed here, the interest is in quantifying the cost of capital for a hypothetical third party who would assume the remaining risk in the liabilities. The extent to which this party may or may not be subject to these frictional costs is a function of the jurisdiction assumed.

Finally, we note that to the extent included, these frictional costs represent additional costs above and beyond the returns required to compensate capital providers for the risk they assume.

5.2. Impact of Corporate Income Taxes

The discussion above has ignored the impact of corporate taxes, which we think are an important consideration since market-based estimates of the returns required by shareholders are inherently after consideration for corporate income taxes.

We begin by noting that premiums collected reflect both components of the fair value calculation – the present value of expected benefit payments (BEL) and the risk margin (MVM). The tax consequences associated with collecting the BEL in the premiums depends on how and when the liabilities are recognized for tax purposes. An acceptable approximation is to assume that there is no *net* tax consequence associated with that portion of the premium and liabilities. On an expected value basis, the investment income from that portion of the premium is offset by the unwinding of the discount on the BEL.

Of more importance here is the corporate tax treatment of the MVM and the corporate taxes that would be payable on the investment income on the capital. These taxes are important components of the price at which a third party would realistically assume the risk of insurance liabilities and should be reflected.

Three specific assumptions need to be clearly specified. The first is the tax jurisdiction assumed for the hypothetical third party, since effective tax rates vary considerably throughout the world. The second is the extent to which companies need to model the actual tax codes of the relevant jurisdiction or whether more simplistic adjustments to account for corporate taxes can be made. And the third is the extent to which the MVM itself is assumed to be taxed. While we have ignored the timing of income recognition in the present discussion, whether the MVM itself is considered taxable profit at inception or whether it is initially included in the technical provisions (reserves) and treated as taxable income only over time impacts the degree to which additional offsetting tax consequences need to be considered.

Once those details can be assessed, the MVM may need to be grossed up to reflect the effective tax rate on the risk margin. If it is necessary to reflect both the tax on the MVM and the tax on the investment

[14] See Feldblum and Swiss Re studies.

income, it is convenient to simply gross up the *spread* used in the MVM calculation. This can be done using the formulas shown in Appendix G.

5.3. Calibration to Market Prices

Detailed support for any specific annual cost of capital rates to be used in the cost of capital method has not been developed. Using placeholder assumptions, such as the 6% rate used in some regulatory capital applications, is reasonable at this stage of development of MVM standards, since ultimately we expect a pragmatic rather than a theoretical solution to be obtained. This pragmatic solution may simply require all insurers to use a specified rate, perhaps varying by product, without reference to specific models such as CAPM or specific assumptions such as betas or the equity market risk premium.

We note, however, that such broad simplifications may be more acceptable for solvency applications than for financial reporting applications. In the latter case, it may be more appropriate and desirable to use capital costs more specifically based on each insurer's circumstances. This would likely enable individual companies' fair value liability estimates to be more supportable as being consistent with market prices.

This practical approach would, however, place more emphasis on the need to properly calibrate the rate used to provide assurance that the results at least approximately reflect actual market prices. We explore whether this is feasible by reference to market prices for property catastrophe bonds.

5.3.1. Calibration of Property Catastrophe Bond Spreads

The graph below[15] shows the spreads paid on a variety of catastrophe bonds both before and after the 2005 hurricane season.

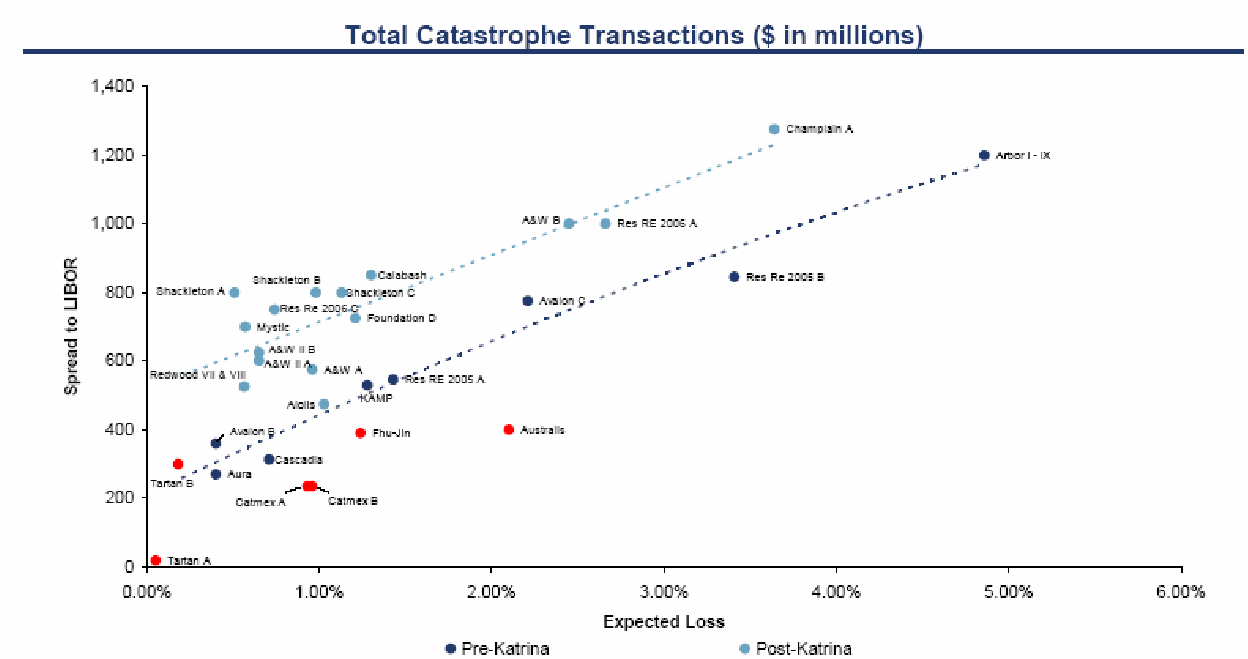

This data can be used to assess the annual cost of capital rate that would produce the same MVM as implied by these market prices.

---

[15] Source: *Presentation Regarding the Transfer of Property Catastrophe Risk to the Capital Markets,* Goldman Sachs, August 26, 2006.

As in any empirical analysis such as this approximations must be used. Here, we do not have detailed information regarding the claim cost distribution for each of the bonds shown. However, we used a "typical" loss distribution reflecting a diversified property-catastrophe risk portfolio as a baseline property-catastrophe distribution. This distribution, which is assumed to have a maximum possible claim size of $500 million, is shown below:

Table 15: Assumed Property-Catastrophe Ground-Up Claim Distribution

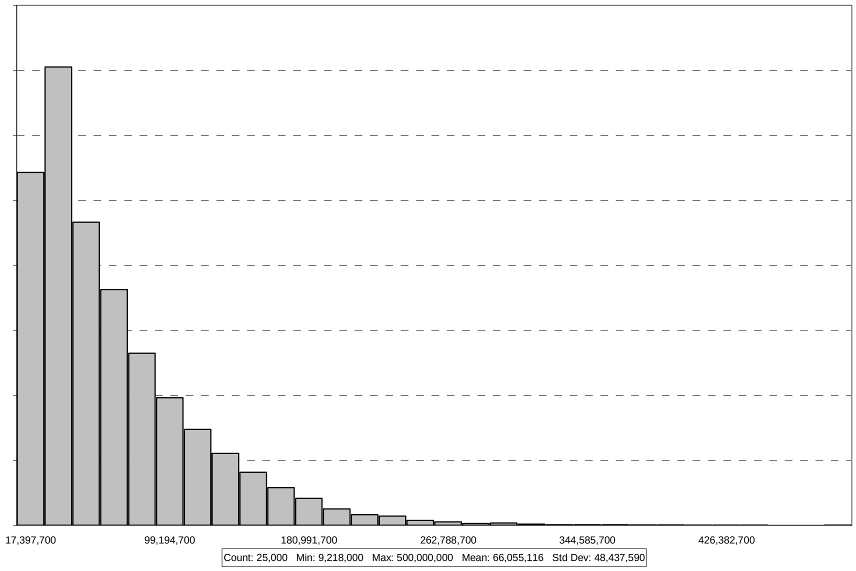

This distribution has expected claim costs of approximately $66 million or 13.2% of the maximum limit of $500 million. To compare the pricing indicated for a range of transactions with different expected loss amounts relative to the maximum possible loss (the notional amount of the bond), we varied the attachment point of this generic distribution until the expected loss ranged from 2% to 6%.

For each case, we assumed the risk period was a single year and that the bond matured in 1.5 years to allow for some settlement lag[16]. The following table shows the implied spread (ignoring tax considerations) that would have to be used within the cost of capital method to match the MVM implied by prices indicated by the line of best fit in the graph above.

[16] Note that both of these assumptions are significant simplifications. Most of the property-catastrophe bonds shown in the graph are likely to cover multiple risk periods and have an overall maturity of 3 or more years. However, we do not believe that this materially affects the conclusions drawn here.

Table 16: Implied Cost of Capital Rate Based on Market Price for Catastrophe Bonds

Max Loss:	500,000,000	500,000,000	500,000,000
Attachment Point:	105,934,292	69,323,626	46,849,292
Policy Limit:	394,065,708	430,676,374	453,150,708
Mean Loss in Layer	7,881,059	17,227,434	27,189,273
Expected Loss as % of Limit:	2.00%	4.00%	6.00%
99.5th Percentile in Layer	152,127,283	188,737,949	211,212,283
Capital Amount:	144,246,224	171,510,516	184,023,010
Pre-Katrina Spreads			
Cat Bond Spread at % of Limit	6.50%	11.00%	15.50%
Cat Bond Premium, Spread X Limit:	25,614,271	47,374,401	70,238,360
Cat Bond Risk Margin:	17,733,212	30,146,967	43,049,087
Risk Free Rate (assumed):	4.00%	4.00%	4.00%
Implied Cost of Capital:	**8.58%**	**12.27%**	**16.33%**
Post-Katrina Spreads			
Cat Bond Spread at % of Limit	9.00%	13.00%	17.00%
Cat Bond Premium, Spread X Limit:	35,465,914	55,987,929	77,035,620
Cat Bond Risk Margin:	27,584,854	38,760,495	49,846,347
Risk Free Rate (assumed):	4.00%	4.00%	4.00%
Implied Cost of Capital:	**13.35%**	**15.77%**	**18.90%**

Three important observations can be drawn from this example.

- The implied cost of capital rate varies for each sample transaction. This calls into doubt the validity of a single cost of capital rate for all insurance risks;
- The implied cost of capital rate changed significantly after the 2005 hurricane season. This suggests that the rate used would have to be routinely recalibrated in order for the resulting MVM to remain meaningful.
- The implied (post-Katrina) spreads are significantly in excess of the placeholder rates (6%) now being used in certain solvency applications of this methodology.

5.3.2. Observable Rates Are Blended Rates for Hedgeable and Non-Hedgeable Risks

A final issue with regard to calibration is the fact that the MVM framework is intended to be applied only to non-hedgeable risks. The reason for the emphasis on non-hedgeable risks only is that it is assumed that that risk margins for hedgeable risks are adequately captured in the market prices of assets and in the valuation of the BEL for insurance risks (see Section 2.2). This raises a question as to whether it is possible to directly calibrate the cost of capital for non-hedgeable risk separately from the hedgeable risks using market rates of return on invested capital.

Consider the case of two insurers with significantly different investment portfolio strategies but the same current portfolio value. One insurer invests solely in short term US Treasury securities; the other insurer invests solely in equities.

If neither company had any insurance liabilities, then we would expect the rates of return required by their respective shareholders to vary but the *value* of their investment portfolios would be the same. Once insurance liabilities are assumed to exist, the risk profile of both firms would change and the *overall* compensation for risk demanded by their respective shareholders would also change. Even if the insurance risks are identical, the required rates of return would still differ in part because they would reflect the differential risk profile of the assets.

Therefore, if we intend to strip out only the non-hedgeable risks and try to quantify the rate of return required for assuming just those risks, it appears unlikely that the market data for insurance companies' equity returns will be useful. These data reflect a blend of capital costs for the asset risks and the insurance risks (i.e. the hedgeable and non-hedgeable risks).

6. Considerations for Solvency Applications

6.1. Introduction

Both the fair value financial reporting and solvency applications require the calculation of the fair value of the liabilities, comprised of the BEL and the MVM. In previous sections of this report, we have focused solely on the financial reporting applications and highlighted specific implementation issues associated with the cost of capital method. These issues remain relevant when applying the same general method in solvency applications. In addition, there are other issues that are specifically relevant to a solvency application.

This section of the report will present a broad solvency framework that uses fair value as the basis for measuring insurance liabilities and then address implementation challenges associated with the cost of capital method.

6.2. Capital Standards for Insurer Solvency

Several insurance regulatory bodies in the EU, Switzerland and elsewhere are working to revise and harmonize capital requirements for the purposes of insurer solvency regulations. A common theme among these efforts is to extend the use of fair value concepts being adopted for financial reporting purposes to also provide a framework for establishing minimum standards for the capital account.

Specifically, in addition to reporting the technical provisions as of the valuation date on a fair value basis, minimum capital requirements related to non-hedgeable insurance risks must be sufficient to absorb the impact of *changes* in the fair value of the insurance liabilities over a specified solvency time horizon, such as one year.

In these frameworks, the goal is to ensure that the initial capital is at least as great as the potential change in the fair value under "severely adverse conditions". Although the definitions vary, one common approach is to define these adverse conditions as those represented by the 99.5^{th} percentile of the range of possible fair values at the end of the selected solvency time horizon. These adverse conditions are thought to represent a *distress scenario for the company*. The objective of the capital standard is to ensure that, under such a distress scenario, the company could restate its liabilities to their fair value and still have a positive balance in its capital account[17].

To achieve this objective, it is necessary to quantify the following components:

- <u>Current Best Estimate Liability (BEL_0)</u>

  This component reflects the *current* BEL and is identical to the quantity discussed in the financial reporting context.

- <u>Change in BEL</u>

  This component is also referred to as the *Solvency Capital Requirement for Non-Hedgeable Risks* under Solvency II. As we mention briefly in Footnote 17, the total solvency capital requirement may include provisions for other risks such as market, credit or operational risk as well. We will nonetheless refer to the solvency capital requirement for non-hedgeable (insurance) risks as the "SCR" throughout this report for convenience.

[17] It is important to note that this describes the capital requirements associated with the insurance risks only. Additional capital requirements may also exist for market risk, credit risk and operational risks as well. However, throughout this report our focus is on the calculation of the MVM for insurance risks. We therefore do not address these other components of solvency capital standards or other aspects of solvency regulations.

The SCR (for non-hedgeable risks) reflects the potential change in the BEL over a one-year time horizon[18]. Given a full distribution of possible values for the change over the time horizon, the SCR represents the change in value at a reasonably high percentile, such as the 99.5th percentile.

The sum of the current BEL and the potential change in the BEL is merely the *potential* BEL at the end of the horizon, assuming that the 99.5th percentile events have occurred. This can be thought of as a *distress scenario* BEL, which we denote as BEL^{DS}_{1}.

These other components of the total solvency capital requirements, including those related to potential changes in the BEL for hedgeable insurance risks, will be ignored so that we can remain focused on the issues associated with the use of the cost of capital method for measuring the MVM for non-hedgeable insurance risks.

- Current Market Value Margin (MVM_0)

  This component reflects the *current* MVM and is identical to the quantity discussed in the financial reporting context.

- Change in Market Value Margin

  This component reflects the potential change in the market value margin during the (one-year) solvency time horizon.

  The sum of the current MVM and the change in the MVM could be combined to reflect the MVM at the end of the time horizon, or MVM^{DS}_{1}. In most cases it would be mechanically easier to quantify this aggregate amount, MVM^{DS}_{1}, rather than its components.

  Nonetheless, the solvency application calculations are simply fair value calculations as described in the previous section with two significant differences – they are done as of the *end* of the time horizon and they are contingent upon a distress scenario for the company.

  The difference between MVM^{DS}_{1} and the initial MVM_0 will reflect the combined effect of each of the following components:

  - Market Price per Unit of Risk Changes – During the course of the one-year period, the price (per unit of risk) of transferring the risk associated with the liabilities could change.
  - Perceived Risk Changes in a Distress Scenario – During the course of the one-year period, the perceived amount of risk could change in response to observed claim and benefit activity during the year. In some cases, actual experience will deviate sufficiently from what was expected (i.e. a distress scenario) that underlying risk models may be altered. The market's recalibration of hurricane risk models after the 2005 hurricane season is a good example of this.

    In other cases, the actual experience, even when adverse, may be interpreted as resulting from random variation and the risk of further adverse deviations may actually decrease.

    It is difficult to know, *ex ante*, precisely which effect will dominate.
  - Passage of Time Implies Less Remaining Risk – All else being equal[19], the MVM in one year will be lower than at inception simply because of the passage of time. If the MVM

[18] For clarity, here and throughout this discussion we assume that this distress scenario BEL is evaluated at the end of the period but just prior to any claim or benefit payments for the period are made. This allows us to avoid complicating the definition with references to the actual or expected claim and benefit payments made during the period.

[19] Specifically, all else being equal refers to the absence of i) changes in the market price per unit of risk, ii) the perceived risk and iii) the total risk exposures. In the case of the latter component, we mean that no additional insurance policy risk other than what was included in the initial calculations exists at the end of the time horizon.

is calculated as the cost of maintaining capital to support the risk in the liabilities over all future periods, then MVM_1 differs from MVM_0 in an obvious way – it simply ignores the cost of capital contribution from the first period. This is shown in the graph below which uses all of the periods' values for calculating MVM_0 for financial reporting purposes and uses only the period 2 through period 8 values for calculating MVM_1 for solvency purposes.

Table 17: Comparison of Initial MVM and End of Period MVM

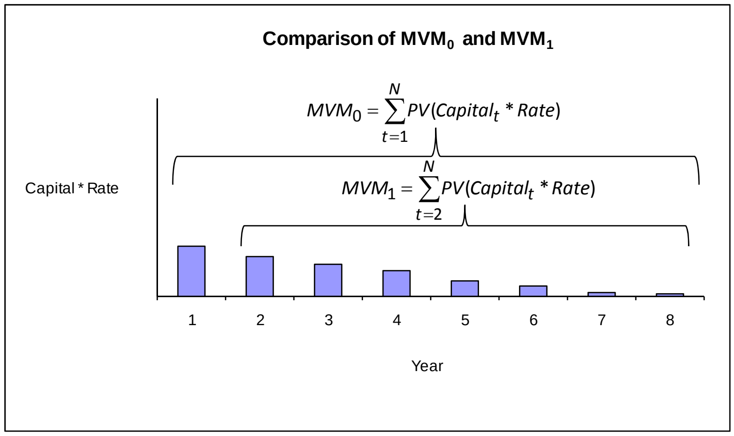

Each of these components could cause the MVM to change over a one year horizon. In our numerical examples in this report we will primarily reflect the final component only. This makes the calculation of the risk margin for solvency purposes trivial once the risk margin for financial reporting purposes is calculated. In practice though, significant consideration should also be given to the need to incorporate the potential for changes in the market price per unit of risk and the potential changes in the perceived risk, both of which could be significant in a distress scenario. This is because the total "required capital" must be sufficient to absorb changes in both the BEL and changes in the risk margin.

This point is emphasized by reference to the following graphical depiction of the relationship between the components of the solvency balance sheets at the valuation date and one-year forward:

Table 18: Solvency Balance Sheet Components (Insurance Risks Only)

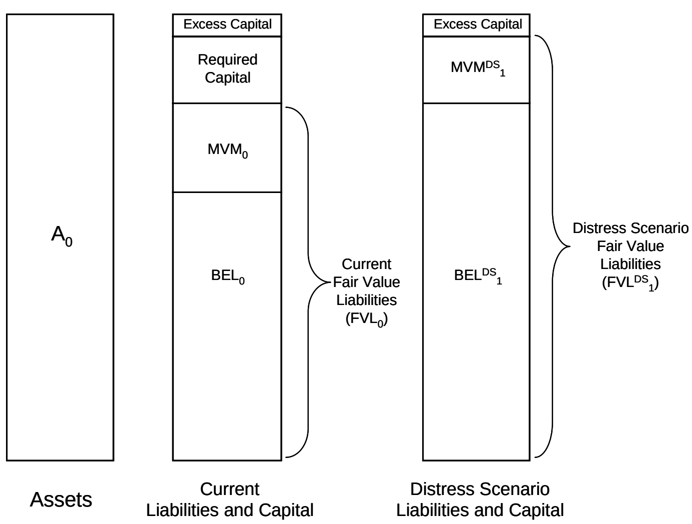

6.2.1. *Required Capital and the Change in Fair Value*

In the diagram above, the "required capital" reflects the potential change in the fair value liability. We used the "DS" superscript to indicate distress scenario values and we used the "0" and "1" subscripts to denote values at the valuation date (time t = 0) and one year hence (time t = 1), respectively.

As described above, the required capital must account for both the potential change in the BEL and the potential change in the MVM over the selected time horizon.

$$\begin{aligned}\text{Potential Change in Fair Value} &= FVL_1^{DS} - FVL_0 \\ &= (BEL_1^{DS} + MVM_1^{DS}) - (BEL_0 + MVM_0) \\ &= (BEL_1^{DS} - BEL_0) + (MVM_1^{DS} - MVM_0) \\ &= \Delta BEL + \Delta MVM\end{aligned}$$

The magnitude of the change in the BEL, denoted ΔBEL above, and the change in the MVM, denoted ΔMVM above, depends on the time horizon selected and the degree of confidence desired. For instance, under Solvency II and the Swiss Solvency Test, the intent is to set capital requirements for insurers such that there is a high probability, such as 99.5% under Solvency II, that at the end of a **one-year time horizon** the firm's assets will be no less than the sum of the distress scenario BEL and the distress scenario MVM, both as measured at the one-year horizon period.

6.2.2. *Change in Net Assets*

The framework described above can also be depicted in terms of the change in the firm's net assets. This alternative definition would also incorporate potential changes in the assets that result from market risk, credit risk or operational risk. We depicted the required capital as being driven by the potential change in the fair value of the liabilities and ignored these other components here to focus on the capital required to support non-hedgeable insurance risks and the calculation of the MVM.

6.3. Implementation of Cost of Capital Method in Solvency Applications

In the solvency application, the goal is to ensure that the capital held at the beginning of the period is at least sufficient to cover the *change* in the BEL during the selected time horizon and the *change* in the MVM, also during this horizon[20]. These two components comprise the *change* in the fair value liability, creating the need to specify each of the following:

- Risk Exposure Horizon Used to Estimate the Change in the BEL
- Risk Exposure Horizon Used to Estimate the End of Period MVM
- Appropriate Assumptions for Calculating the MVM in a Distress Scenario
- Appropriate Assumptions for Calculating the End of Period MVM

6.3.1. *Risk Exposure Horizon Used to Estimate the Change in the BEL*

While not directly related to the cost of capital methodology for calculating the MVM, an important component of the solvency capital requirements is the change in the BEL during the solvency time horizon. For instance, under Solvency II the solvency capital requirement for non-hedgeable insurance

[20] Additional capital requirements for market risks, credit risks and operational risks may also exist. We limit our discussion to capital requirements for non-hedgeable insurance risks and therefore refer to the changes in the fair value of the insurance liabilities as the "sole" determinants of the capital requirements.

risks (SCR) reflects the 99.5th percentile amount by which the BEL could change during a one-year period.

In Section 3.3 we presented two alternative perspectives that could be used to quantify the MVM, one of which assumed that the capital to support the ultimate risk had to be fully funded at inception and one that assumed that the capital could be funded sequentially using a series of one-year risk exposure horizons. We pointed out that either perspective could theoretically be used to calculate the MVM, but that the choice of annual cost of capital rates would have to differ depending on which perspective was used. Further, we argued that calibration to actual market prices may be easier using an ultimate risk exposure horizon because of its comparability with the capital base commonly used by insurers and reinsurers to price transactions using existing regulatory and rating agency capital models.

Similarly, the same risk models used to quantify the capital base over these two alternative risk exposure horizons can be used to quantify the change in the BEL. But for purposes of determining the change in the BEL over a one year horizon, only the one-year, sequential approach is appropriate.

To see why this is the case, note that when the change in the BEL is added to the beginning BEL (BEL_0), this produces the end-of-period BEL (BEL^{DS}_1). This quantity BEL^{DS}_1 reflects a *distress scenario* best-estimate liability assuming that the risks that can materialize within the one-year period have occurred. For long term insurance liabilities considerable risks may still remain at the one-year time horizon, but we are interested only in the *expected* value of the claim and benefit payments at that point. In this sense, BEL^{DS}_1 is conditional upon the first year's adverse experience but remains an *expected value* with respect to future cash flows, albeit an expected value that is conditional on the occurrence of the events that gave rise to the distress scenario. It is sufficient to absorb the effect of only these events that occur during the one-year period.

If instead the quantity BEL^{DS}_1 were sufficient to cover all future adverse conditions rather than just the adverse experience over one year (and conditional expected values thereafter), it would itself adequately fund these adverse events. Additional risk margins would not be needed to supplement these funds.

6.3.2. Risk Exposure Horizon Used to Estimate the End of Period MVM

The previous section concluded that a one-year risk exposure horizon is appropriate to quantify the change in the BEL. This is true regardless of which of the two perspectives discussed in Section 3.3 is used to quantify the distress scenario MVM – either the ultimate risk exposure horizon or the one-year risk exposure horizon.

Recall from Section 3.3 that the two alternative perspectives for calculating the MVM produced the following capital base in each period for a 5-year term insurance portfolio:

Table 19: Comparison of Capital Base Using Ultimate and One-Year Risk Exposure Horizons

	Capital Base for MVM Calculation	
Year	Ultimate Risk Measure	One-Year Risk Measure
1	1,827,604	913,655
2	1,665,323	910,274
3	1,470,311	907,705
4	1,254,682	905,970
5	907,415	905,091

Depending upon the annual cost of capital rate used, either of these series of risk measures could be used to define the capital base each period and to estimate the MVM. If the ultimate risk exposure horizon were used for this purpose, then the capital base used to quantify the distress scenario MVM would be the amounts that are shaded in the middle column and labeled "Ultimate Risk Measure". Even when this is done though, the change in the BEL must reflect the amount shown shaded in the right-most column.

That is, even when the MVM reflects an ultimate risk exposure horizon, the SCR must reflect a one-year risk exposure horizon.

For example, as demonstrated in Section 3.3.4, the BEL at inception is $8.34 million and the MVM using an ultimate risk exposure horizon and term-specific percentiles is $324,445, at inception, for a total fair value of $8.66M.

Subject to a distress scenario, and immediately prior to claim and benefit payments being paid, the best estimate liability could increase by $913,655. The MVM at this point[21] would then reflect the highlighted capital base shown above in Table 19, as follows under the assumption that the perceived risk is unchanged and that the cost per unit of risk is also unchanged:

Table 20: Calculation of Distress Scenario MVM – Ultimate Risk Exposure Horizon

Year	Capital Base	Annual Cost	PV Annual Cost
2	1,371,465	82,288	78,369
3	1,295,476	77,729	70,502
4	1,195,622	71,737	61,969
5	907,415	54,445	44,792
		MVM:	**255,633**

Here, the distress scenario "fair value" would reflect the potential BEL^{DS}_1 = $8.34M + $913,655 = $9.25M and the potential MVM^{DS}_1 = $255,633 for a total fair value of approximately $9.51M. This is a change of approximately $9.51M - $8.66M = $845,000. This is the total amount of capital that would have to exist on the company's balance sheet, at inception, so that there would be assurance that the liability could be reported at fair value in a distress scenario.

Notice that unless specific assumptions are made regarding how the perceived risk may change as the result of a distress scenario and how the cost of capital rate may change as well, the change in the MVM will *reduce* the overall capital requirements reflected in the SCR. This is seen in the numerical example above where the total change in the fair value liability, including the change in the MVM, is $845,000 while the SCR is $913,655.

As an alternative, the same calculations could be performed using sequential, one-year risk exposure horizons. This approach is used, for instance, under the Swiss Solvency Test and the proposed Solvency II framework. There, the same one-year risk exposure horizon calculations are used for both the SCR calculation and for the MVM calculations, reflecting the values in the last column of Table 19. Using those amounts, the only difference with the calculation above is that the MVM would reflect the alternative series of capital amounts in both the initial calculations and the distress scenario calculations. For instance, the distress scenario MVM would be as follows:

Table 21: Calculation of Distress Scenario MVM – One-Year Risk Exposure Horizon

Year	Capital Base	Annual Cost	PV Annual Cost
2	910,274	54,616	52,016
3	907,705	54,462	49,399
4	905,970	54,358	46,957
5	905,091	54,305	44,677
		MVM:	**193,048**

[21] Note that for convenience we have not recalculated the required capital conditional on the distress scenario occurring. As a result, the required capital amounts would in practice differ slightly from those shown here, which assume that experience to date has been as expected rather than adverse.

Here, because the same 6% rate is used for the annual cost of capital, the resulting MVM is lower than in the previous calculation. But otherwise, the BEL and SCR components are the same. Clearly one would need to be careful to properly adjust the annual rate used so that it is appropriate for the choice of risk exposure horizons.

In addition, note that this example used simulated sequential calculations of the one-year risk measure. In practical implementation of this methodology (e.g. under the Solvency II QIS 3 field testing), it is generally assumed that the initial one-year risk measure (the SCR shown above) can be used along with the future outstanding BEL, as a ratio to the initial BEL, to determine the capital base in each period. If this were done here, the following alternative calculations would have resulted:

Table 22: Calculation of Distress Scenario MVM – One-Year Risk Exposure Horizon (Approximate)

Year	BEL	Ratio to Initial BEL	Capital Base	Annual Cost	PV Annual Cost
1	8,338,097	100%	913,655	0	0
2	6,755,724	81%	740,265	44,416	42,301
3	5,134,530	62%	562,621	33,757	30,619
4	3,467,711	42%	379,978	22,799	19,694
5	1,757,760	21%	192,608	11,556	9,508
				MVM:	**102,122**

In practice, care would have to be taken to ensure that this approximation of the future periods' capital base accurately reflected the risks in each period. In this particular case this was not a reasonable approximation.

6.3.3. Appropriate Assumptions for Calculating the MVM in a Distress Scenario

In the solvency application of the cost of capital methodology, the intent is to capture the risk margin that would be appropriate *after* a distress event. The assumption is that the company holds enough capital to absorb the impact of an extreme change in the BEL over a single period and then, after this extreme event still has sufficient capital to pay the appropriate MVM to a third party so that they will assume the risk associated with the (revised) liability.

This means that the MVM for use in the solvency context, denoted here as MVM^{DS}_1, must not only reflect the market environment that will exist in one year, but it must also be conditional upon the extreme change in the best estimate liability occurring during that period. This was ignored in the simplified calculations shown in the previous section, which relied upon risk measures calculated at inception. In practice, ignoring this aspect of the calculation may be inappropriate. However, making these adjustments accurately may be fairly complicated.

For instance, consider a multi-year P&C insurance policy covering loss events over a two-year period but containing an aggregate policy limit. In the absence of any loss events occurring in the first period, the aggregate limit remaining for the second period will be the same. Hence, the risk margin may still be considerable and perhaps comparable to the initial risk margin. However, if the distress scenario event were to occur in the first period, the remaining policy limit and remaining risk may be substantially reduced, perhaps to zero. In this case, the MVM^{DS}_1 would be substantially lower than at inception.

An offsetting effect is that when a distress scenario event occurs in the first period, market participants may reevaluate their risk models. This will impact the value of the distress scenario risk margin, MVM^{DS}_1. This effect, which could be quite significant, may be impossible to quantify reliably on an *ex-ante* basis. The confidence that the company will have adequate resources to transfer the risk associated liabilities in a distress scenario may therefore in practice be lower than is intended by various solvency frameworks.

In either case, the calculation of the *conditional* MVM^{DS}_1 may involve more complex calculations than a simple extension of MVM calculations performed for fair value financial reporting. To stay true to the intent, the calculation is not as simple as shown earlier, where we relied on the calculations performed for the fair value financial reporting calculations and then merely eliminated the first period's capital costs.

6.3.4. Appropriate Assumptions for Calculating the End of Period MVM

When used in the financial reporting context, the quantity of interest is the MVM that is applicable at the valuation date, which can be assessed using information available at the time. In the solvency context the quantity of interest is the MVM that we believe will apply at the *end* of the horizon.

In the previous section we noted that when the capital base is measured to reflect the amount conditional on a distress scenario, the capital base for the remaining risk may be relatively difficult to quantify. In addition, the annual rate of return applied to this capital amount under the cost of capital method may also differ at the *end* of the period.

The MVM at the end of the period, by definition, cannot be calibrated properly using current market data unless the appropriate cost of capital rate remains constant. If instead the intent is to reflect the potential for this rate to change at the end of the period, it may not be appropriate to use the same cost of capital rate used for fair value financial reporting purposes.

7. Benchmarking Results

7.1. Introduction

In both the financial reporting and solvency applications, the *intent* is to ensure that the sum of the best estimate liability and the MVM adequately reflect the price an insurer would have to pay to transfer the remaining risk in the liability to a third party.

Such fair value measurements are easier to obtain when valuing financial assets because the models used can be calibrated to match actual market prices for other assets and marketable securities with identical or nearly identical cash flows. For most insurance liabilities, the lack of market price data and the unique nature of insurance liabilities make this process far more difficult.

As a result, it is important to evaluate the degree to which the cost of capital method can achieve the desired level of reliability and consistency. In this regard, there are several issues to consider:

- Using Entry Prices to Calibrate Exit Prices
- Availability and Granularity of Benchmark Prices
- Challenges Associated with Capturing Current Market Conditions

7.2. Using Entry Prices to Calibrate Exit Prices

There has been considerable debate regarding the merits of using entry prices, as indicated by the premiums actually charged by the insurer when issuing the policies, as the basis for future valuations of the liabilities. This debate relates primarily to the timing of the recognition of income over the life of the liability. While we consider this to be an important issue, for the present purposes we are more concerned with the usefulness and reliability of the MVM method to produce the proper value and do not address income recognition issues here.

With respect to the valuation issue, we recognize that entry prices can offer an objective assessment of an actual transaction price as it reflects the price that two parties did in fact agree to pay and receive in an actual transaction. However, entry prices are limited as a mechanism for evaluating the price at which the risk associated with the liability could be transferred to a third party on future dates as the result of the following:

- Impact of Portfolio Effects and Diversification – Entry prices often reflect the value of the policy to an individual *insured*. The sum of the prices for each policy is likely to overstate the price at which an insurer could combine all of its gross exposures into a single portfolio and transfer the risk associated with the entire portfolio to a third party.
- Change in Valuation Over Time – By definition entry prices cannot capture the effects of changes in the market's perception of the risk or changes in the market price for that risk, including those that result from insurance pricing cycles. Using entry prices provides no assurance that the reported valuations at any time, other than at inception, accurately reflect the cost that would be incurred to transfer the risk associated with the liability to a third party.

As a result of these concerns, it is argued that market value margins (and the associated best estimate liabilities) are best reported by reference to *current* assessments of aggregate risk and the market price of risk that prevails on the valuation date. We agree with this conceptually, but note that exit prices are only meaningful to the extent that each of the following is true:

- they can be calculated consistently across different insurers and across different products;
- they can be calibrated against objective benchmarks that reflect actual prices at which the risks associated with insurance liabilities are transferred in arms-length transactions; and

- they rely upon parameters and assumptions that can capture both the *current* market environment for financial reporting purposes as well as capture the potential *future* market environment in a distress scenario for solvency purposes.

7.3. Availability and Granularity of Benchmark Prices

A focus on exit prices as opposed to entry prices raises the need to be able to calibrate the results of the MVM calculations to objective benchmark prices.

In cases where identical asset or liability cash flows exist in a form that is routinely transferred in arms-length transactions among various parties, such calibration is relatively easy. In cases where the identical cash flow streams don't exist but can be replicated using combinations of other benchmark securities, it may be possible to have reasonable assurance that the prices obtained are calibrated to match actual prices that would be paid. Many classes of insurance liabilities do not meet either of these two criteria.

Insurance liability cash flows are highly uncertain as to both timing and amount. This uncertainty alone does not preclude the ability to use benchmark securities to value the liability cash flows. For example, some insurance products have cash flows that can be largely determined with reference to underlying risk drivers such as equity indices or interest rates. In these cases, the prices of marketable securities can be used to establish a risk-adjusted value for the liabilities (e.g. through the use of replicating portfolios). But for many insurance products the uncertainty in claim frequency rates, claim severities, policyholder lapse rates or policyholder option election rates are tied only loosely, if at all, to observable or hedgeable risks. For these products, prices of marketable securities are not sufficient to value these risks. Instead, specific insurance product and price data is needed.

The following options can be considered:

- Primary or Reinsurance Market Premiums – It is sometimes argued that because insurance liabilities are not traded on an organized, liquid and open market it is not possible to reliably obtain "market" prices for these liabilities. While these facts significantly impair the precision of any such estimates, we believe this argument is somewhat overstated since these limitations apply to many thinly traded financial instruments as well.

  Nonetheless, several limitations exist, including the following:

  - Diversity – The diversity of product features and risk characteristics that exist for many insurance products means that there are fewer "generic" benchmark transactions that can be used to replicate the risk characteristics of specific liabilities.
  - Transparency – The lack of published or readily available price indices for insurance and reinsurance transactions significantly limits the ability to use these transactions as a guide to calibrating market value margins.
  - Aggregation of BEL and MVM – In many market transactions where the prices for the transfer of insurance risk can be observed, it is not possible to disaggregate the BEL and the MVM. As a result, any attempt to calibrate the calculation of the MVM is contingent upon properly calculating the BEL as well. Using the wrong estimate of the BEL could imply a misleading MVM.
  - Magnitude – Unlike in the case of a financial institution where the valuation of thinly traded and illiquid assets may comprise an insignificant portion of the fair value balance sheet, thinly traded and illiquid insurance liabilities dominate the balance sheets of insurers. This makes concerns about the precision of these estimates perhaps more significant in the aggregate.

- Loss Portfolio Transfers and Closed Block Transactions – In certain cases, prices for loss portfolio transfers and closed block transactions are publicly reported and could potentially be used to infer the magnitude of the risk margins embedded in the transactions. However, inferring the magnitude of the risk margin would also require knowledge of each party's specific assumptions regarding the best estimate liability, which is unlikely to ever be disclosed.

  In addition, it must be recognized that often there are a host of related and unrelated motivations for executing these transactions. Unlike marketable securities whose prices reflect only the cash flows and risks of the specific securities included in the transaction, prices for loss portfolio transfers and closed block transactions often reflect a variety of strategic objectives for one or both parties. Inferring the "price of risk" alone may not be possible from such complex transactions.

- M&A Transaction Prices – The most reliable estimate of the valuation of insurance liabilities that could be obtained in arms-length transactions is available from insurance M&A transactions. There are, however, two specific limitations to using this data for the purposes discussed here.

  The first limitation is that these transactions involve diversified portfolios of insurance risks. The aggregate risk margin reflects the degree of diversification inherent in the portfolio. Further, separation and allocation of the risk margins to specific products of lines of business may not be possible.

  The second limitation is that, except in limited cases where a company or its insurance portfolio is purchased solely for the purposes of placing it into runoff, a significant component of the price paid reflects the company's *franchise value*. This franchise value reflects the risk-adjusted value of *future* business that is not relevant in either the financial reporting or solvency applications addressed here.

7.4. Challenges Associated with Capturing Current Market Conditions

Estimates of the fair value of liabilities for financial reporting purposes must always reflect current market conditions with respect to risk assessment and the market price of risk. This will create challenges for regulators and companies to ensure that this can be achieved in a manner that is both practical to implement and useful to third parties.

- Practicality – Given the regulatory and oversight contexts in which the market value margins are used, it may be impractical for regulators and industry groups to regularly monitor and modify the prescribed parameters or to review and approve 'principles based' judgments about the parameters used by management.
- Usability – A primary motivation for adopting market-based financial reporting and solvency standards is to make reported financial statements more useful to third parties, including policyholders, investors and regulators. If market value margins fluctuate significantly from period to period for reasons that are unrelated to companies' assessments of the risks or to the market price of risk, reported results at any point in time could prove to be less useful rather than more useful relative to existing financial statements.

As a separate consideration in the solvency context, the methodology, parameters and assumptions must be capable of capturing the market value margins that could possibly exist at the end of a one-year reporting horizon following a significantly adverse event. Unlike in the financial reporting context in which it is the *current* market conditions that are relevant, the solvency context requires a *forward-looking* assessment of these market conditions. This makes calibration of the key assumptions used for this purpose particularly difficult.

One example of the instability of the market price for risk after an adverse event and the challenges associated with looking forward one period for solvency calculations of the MVM can be found in the property-catastrophe bond market before and after the 2005 hurricane season. We see that price changes were significant after these events, even after correcting for changes in the market's estimates of the best estimate liability, and that estimates of the market price of risk at any point in time may understate the market price of risk in a distress scenario.

The graph below[22] shows the premiums (spreads) charged for a variety of catastrophe bonds:

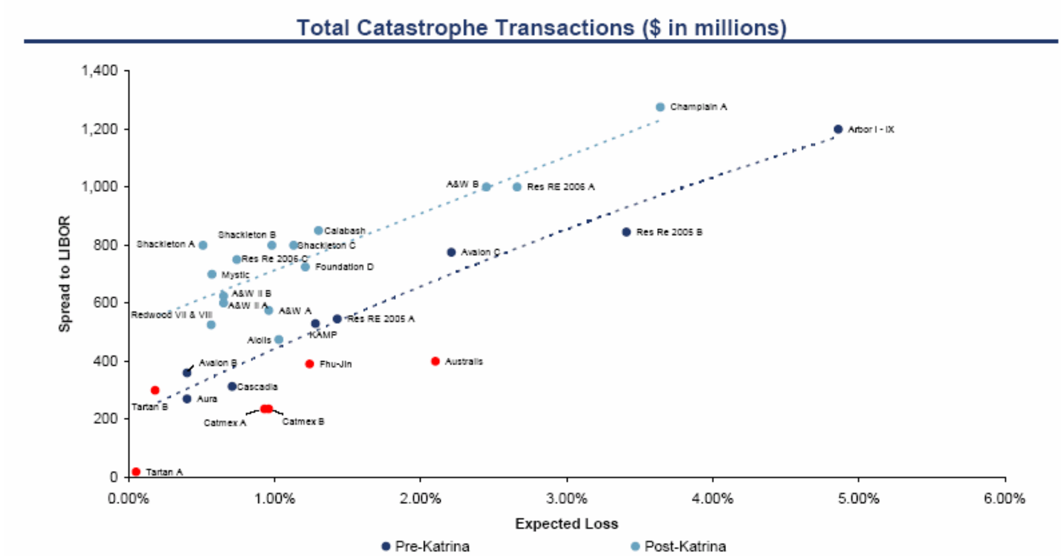

As shown in the graph, the various bonds exhibited a relatively wide range of expected losses and there appears to be a relatively linear relationship between the expected loss and the premiums (spread) charged. The difference in the spreads and the expected loss approximates the market value margin. For instance, the Avalon C bond had an approximately 2.2% expected loss and a premium of approximately 7.0%, suggesting a market value margin of 4.8%. Other bonds show market value margins that are larger or smaller multiples of the expected loss, indicating the importance of proper evaluation of expected losses and risk when quantifying MVM's.

The interesting and relevant observation from this graph though is the fact that after the 2005 hurricane season, which saw significant losses from hurricanes Katrina, Rita and Wilma, the price of risk rose substantially. For instance, post-Katrina the A&W B bond with expected losses of roughly 2.5% (after revisions to the catastrophe models used to estimate the expected loss) had a market value margin of 7.5%. This reflected a significant increase from the 4.8% market value margin for the Avalon C bond that had a comparable expected loss prior to Katrina. This increase reflected a combination of a change in the market's perception of the magnitude of the risk as well as a change in the market's cost per unit of risk.

Appendix H contains similar evidence of the instability of risk margins in S&P 500 index options and credit default swaps (CDS).

[22] Source: *Presentation Regarding the Transfer of Property Catastrophe Risk to the Capital Markets,* Goldman Sachs, August 26, 2006.

8. Appendix A: Percentile Method

8.1. Introduction

An alternative method for estimating market value margins, referred to as the *percentile method*, was originally proposed for Solvency II purposes but has since lost support relative to the cost of capital methodology. We briefly review the merits of this method, paying particular attention to whether or not this approach is more or less practical, reliable or robust relative to the cost of capital methodology.

8.2. Description of Percentile Method

The percentile method uses the underlying risk distribution to directly determine an aggregate fair value of the liability (i.e. the sum of the BEL and the MVM) as a specified percentile of the distribution. Rather than calculating an explicit MVM as in the cost of capital methodology, the percentile method simply selects a specific percentile of the risk distribution as the fair value of the liability. Subtracting the BEL from this amount allows for the determination of the MVM.

In the event that the shape of the risk distribution is stable from product to product and all other underlying assumptions are identical, then the same percentile could be used across all products. However, any differences in the shape of the distribution, correlation with the rest of the risk portfolio, the time horizon over which capital is committed or the annual cost of capital rate, would have to be reflected using a different percentile.

A variation of this method uses a combination of various percentiles instead of a specific percentile. For example, the commonly used Conditional Tail Expectation (CTE) measures the average value of all percentiles above a specified percentile. While mechanically different, this methodology is considered to be conceptually equivalent to the percentile method and will not be addressed in detail below.

8.3. Comparison to Cost of Capital Method

While no *single* percentile is likely to be appropriate for all products or all companies, there is nothing inherently inconsistent with using an explicit MVM calculation using the cost of capital method or determining the MVM using the percentile method. Since the same factors would need to be taken into account in both cases, in practice a table of percentiles appropriate for different classes of business could be produced using the three key inputs to the cost of capital methodology.

The capital base is driven largely by the degree of variability in the underlying risk distribution and so for a wide range of measures of this variability, a range of time horizons over which the capital is held and a given annual cost of capital rate (e.g. 6% per annum here), the following equivalent percentile based MVM estimates can be obtained:

Table 23: Implied Percentile of Loss Distribution from Cost of Capital Method

Equivalent Percentile of Loss Distribution (6% Cost of Capital Rate per Annum)

		Coefficient of Variation		
		0.10	**0.25**	**0.40**
Average Years of Capital Commitment	**1**	59%	63%	66%
	3	71%	75%	79%
	5	81%	85%	88%
	7	88%	91%	93%

As these calculations show, a range of reasonable percentiles useful for calculating the MVM's could be constructed using the same assumptions required for the cost of capital method. Nonetheless, there are some differences in the practical implementation of the percentile method relative to the cost of capital method that should be considered:

- Cost of Capital Method Makes the Underlying Assumptions Explicit – Although the same assumptions are required for both methods, we believe that the need to make those assumptions explicit is a considerable advantage of the cost of capital method and offers valuable opportunities for market participants to assess the calibration of the results.

- Percentile Method Requires Complete Risk Distributions – The percentile method assumes the existence of a complete risk distribution underlying the best estimate liability. Because the same information is, in theory, required to both estimate the best estimate liability and to quantify the initial capital base, this would not appear to reflect any additional burden for implementation. However, in practice this would introduce new challenges because it would preclude the use of many of the approximations intended to be used for the cost of capital methodology, such as the use of scenario analysis or the separate quantification of the capital base for each type of risk prior to calculating an aggregate, diversified capital base.

  When used for solvency purposes the percentile method would also require significantly more complex analysis to properly determine the complete risk distribution one-year forward (and conditional upon the distress scenario).

- Several Challenging Aspects of the Cost of Capital Method are Still Relevant – The numerical example above demonstrates that the appropriate percentile depends upon the time horizon over which the capital is committed and the annual rate that is applied. As a result, the challenges noted previously regarding these two parameters and their proper calibration to actual market prices remain applicable to the percentile method, albeit in a more implicit manner.

- Calibration and Validation Challenges are Not Resolved – All of the challenges associated with calibration of the results discussed for the cost of capital method remain applicable to the percentile method. However, because the cost of capital method shares many commonalities with pricing methods used by many insurers, some aspects of the validation of key assumptions may be easier with the cost of capital method compared to the percentile method.

9. Appendix B: Introduction to Cost of Capital Method

In Section 2, the three key elements of the cost of capital method were summarized as:

- Amount of initial capital base;
- Time horizon for capital commitment;
- Required rate of return on capital per period.

In this Appendix, each of these key elements is demonstrated using a highly simplified numerical example. This example serves only as a basic introduction to the methodology. Each of these elements of the cost of capital methodology is discussed more fully in the main body of the report.

9.1. P&C Liability Application

A simplified numerical example using a stylized P&C liability will set the stage for the more detailed discussion of the cost of capital methodology contained in Sections 3 through 5.

For this example, we assume a P&C insurer with no existing liabilities is established with the intent to write $100 million of premium in a line of business for which the expected present value loss and loss adjustment expense ratio is 70% of premium, up-front underwriting and other operating expenses are 30% of premium and all claims are paid in exactly three years. The capital base is established to ensure, with 99.5% confidence, that claims can be paid if they exceed the 70% expected loss ratio on a present value basis.

We further assume that the following distribution[23] adequately reflects the potential present value claim payments using a fixed risk-free discount rate of 4%:

Table 24: Present Value Loss Distribution – P&C Liability Example

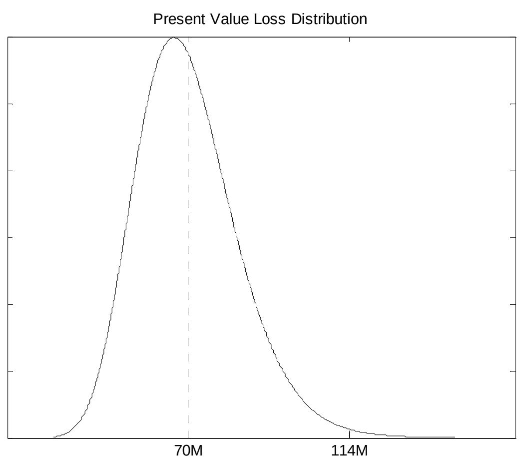

From this picture, the 99.5% confidence standard suggests that we need to have enough capital so that, along with the $70 million of financial resources net of the underwriting and operating expenses paid at inception, we can pay claims with a present value of as much as $114 million. This implies an initial capital base of $44 million. This amount is on a present value basis though, so at the end of each year the capital committed will grow at the risk free rate of return until it is eventually either used to pay claims or returned to the capital providers.

[23] This example assumes the mean loss and loss adjustment expense ratio is 70% on a present value basis. Actual present value loss and loss adjustment expense payments are lognormally distributed with a coefficient of variation equal to 0.20.

Given the simplifying assumption that all claims are paid in three years, we make the related assumption that the initial capital base will need to be held in full for the entire three year period. For the present purposes, we assume that no adjustments are made for the potential release of this capital prior to maturity date of the liabilities, an issue that will be explored in detail later.

The total return per annum required by the capital providers will be assumed, for simplicity, to be 10%. Given the 4% risk-free rate, we can assume that the market value margin need only be sufficient to provide the 6% *spread* per annum.

These three assumptions are combined to produce the following total market value margin, which is shown on a discounted basis to account for the fact that it is collected up-front from the policyholder but paid out to the capital provider, along with the return of their capital, at maturity unless used in whole or in part to pay claims:

Table 25: Calculation of MVM – P&C Liability Example

(1) Period	(2) Beginning Capital Committed	(3) Capital Released	(4) Required Total Return	(5) Assumed Investment Income	(6) Required Spread	(7) Annual Risk Margin	(8) Present Value Factor	(9) Present Value Risk Margin
1	44.32	0.00	10.0%	4.0%	6.0%	2.66	0.962	2.56
2	46.09	0.00	10.0%	4.0%	6.0%	2.77	0.925	2.56
3	47.94	47.94	10.0%	4.0%	6.0%	2.88	0.889	2.56
							Total MVM:	7.67

Notes: (2) = 99.5th Percentile Loss - Expected Loss at inception, adjusted for risk-free investment income thereafter
(6) = (4) - (5)
(7) = (2) * (6)
(8) = [1 + (5)]^-(1)
(9) = (7) * (8)
Total MVM = Sum of (9)

The total MVM of $7.67 million allows, in the case where actual claims are equal to their expected value, for the capital providers to earn a total return of 10% after the investment income on the capital is taken into account.

As this simple numerical example shows, the results depend entirely upon the three key elements noted earlier:

- Amount of initial capital base;
- Time horizon over which that capital must be committed;
- Required rate of return on capital per period.

In each case, the mechanics are trivial but the wide range of actual assumptions that can be used in each stage makes consistent implementation more challenging than it may appear.

Each of these steps is explored further in Sections 3 through 5 and specific implementation challenges are noted.

10. Appendix C: SPDA Example

This appendix uses a simplified Single Premium Deferred Annuity (SPDA) product example to demonstrate the challenges associated with practical calculation of the capital base for MVM calculations in the case of more complex insurance products than were used in the body of the report. In particular, we highlight implementation challenges that arise when products have both non-hedgeable insurance risks as well as hedgeable and non-hedgeable financial risks.

The points raised here are broadly applicable to a wide range of more complex products. The SPDA product is used here, along with other significant simplifications, so that the key points can be presented with minimal references to complications that exist in other products.

10.1. SPDA Product Assumptions

We assume that we have a portfolio of SPDA products with identical product features and an aggregate account value of $100 million. The account balance is adjusted annually at the then-prevailing one-year US Treasury rate as of the beginning of the period *less* 100 bps (1.00%), subject to a minimum crediting rate of 4%.

The product offers the insured the ability, ten years from today, to convert the account balance to a lifetime annuity at then-prevailing market annuity rates. Because we assume that the annuitization rate is not guaranteed (i.e. it will reflect the then-current market rates), we can for simplicity assume that all policyholders receive their account balance at the end of the ten year term, rather than a series of lifetime annuity payments.

The SPDA product typically also offers a death benefit. Reflecting the effect of the death benefits would largely mirror the calculations for the term insurance policy used in Section 3.3. Since repeating the points raised there are not necessary for the present analysis, we will assume that this product does not offer any such death benefits.

Finally, we assume that lapses can occur with no surrender penalties. In practice, declining surrender penalties likely apply to the early years of the policy, but the adjustments to account for such surrender charges are trivial and not required in this example. Therefore, when lapses occur the policyholder is assumed to receive 100% of their account balance and they merely lose the remaining value of the crediting rate guarantee.

10.2. Risk Modeling Components

The risk in this product is driven by the following components:

- Crediting Rate Guarantees – The 4% minimum crediting rate guarantee is the primary source of risk in the product. In the absence of such a guarantee, most other risk can be managed efficiently (see below). However, like all interest rate options, the risk that arises from this guarantee could be hedged if it were not for the additional uncertainty created by unknown mortality rates (if death benefits are included) and lapse rates.

- Lapse Rate Variability – If the lapse rates along each interest rate path were known, the 4% crediting rate guarantee could be hedged efficiently. Even if the lapse rates are not "known", the risk could still be hedged if the relationship between the lapse rates and the interest rates was known. Such dynamic lapse rates that are known functions of the interest rates or other market variables do not materially impact the ability to hedge the interest rate risk. This simplifies the valuation of the liability.

  In practice though, little is known about the precise relationship between interest rates and lapse rates. Even less is known about the potential variability of actual lapse rates from actuarial best estimates of the lapse rates (and/or the dynamic lapse rate functions). The absence of this information prevents a robust stochastic analysis of the value of the lapse risk.

In response to this, a pragmatic approach that has been adopted is to apply "shocks" or "scenario analysis" to baseline lapse assumptions. While practical, this approach raises questions with regard to the ability to properly calibrate these shocks so that they reflect the risk at the desired level of confidence (e.g. the 99.5 percentile). As such, this may introduce unexpected but material differences between the intended and the actual calculations of the risk margins. It may also introduce unexpected but material inconsistencies between the magnitudes of the risk margins for products or risks for which more robust stochastic modeling can be performed (e.g. property-catastrophe risk) and those for which difficult-to-calibrate shocks are applied.

- Investment Strategy – The ability of a firm's invested assets to fully fund its liabilities is impacted by three factors: i) changes in market yields, credit spreads, reinvestment rates, etc. for a given stream of liability cash flows, ii) the actual amount and timing of its liability cash flows for a given set of static market conditions and iii) the combined impact of variability in the cash flows and changes in asset market conditions. These three effects are often referred to collectively as "ALM" risk.

  We ignore these sources of ALM risk in what follows and capture only the risk margin required for the crediting rate guarantee, which is implicit in the calculation of the best estimate liability using risk-neutral valuation methods, and the explicit MVM for the non-hedgeable lapse rate variability. Additional risk margins may be appropriate to account for additional risks introduced as the result of any specific investment strategy.

10.3. Risk Modeling Assumptions

The following specific assumptions are used for the numerical calculations described below:

Product Description

The policyholders have the option to convert their account balance, which grows at a crediting rate net of expenses as described below, to a lifetime annuity at then-prevailing market annuity rates. The option is assumed to be exercisable 10 years from inception.

In-Force

We assume that there are 1,000 policyholders, each with an initial account balance of $100,000.

Crediting Rate

The crediting rate is equal to a reference rate *less* a fixed spread, but subject to a guaranteed minimum crediting rate. The specific parameters are as follows:

- Reference Rate: 1 year US Treasury rate
- Crediting Rate spread: 1.00%
- Guaranteed Minimum Crediting Rate: 4.00%

Lapsation

The lapse rate per annum is the sum of a base rate, a dynamic lapse component that is a function of the spread between the reference rate and the guaranteed minimum rate and a random lapse component.

The base lapse rate is assumed to be 5% per annum.

The dynamic lapse component is equal to a multiple of the spread between the one-year US Treasury rate and the 4% guaranteed minimum crediting rate. To capture the impact of non-hedgeable interest rate sensitivity, we assume that the multiple used is not constant. Instead, the multiple is the absolute value of a standard normal random variable.

The random lapse component is a normal random variable with a mean of zero and a standard deviation of 5%.

Mortality Assumption

For simplicity, the numerical example focuses entirely on the impact of lapsation and the embedded crediting rate guarantee. We have ignored the impact of mortality and assumed that all policyholders will live for the full ten-year term of the option.

At the end of this ten-year term the policyholder will receive an annuity at then-prevailing market rates. We have assumed that there is no additional value impact upon annuitization. Again for simplicity, we assume that the policyholder receives the cash value (their account balance) and do not model the impact of future mortality uncertainty.

Surrender Charge

We have assumed no surrender charges apply.

Expenses

Expenses per annum, as a percentage of account value, are equal to 0.25%.

Stochastic Interest Rate Model

Stochastic interest rate scenarios were generated using a one-factor Hull-White interest rate model. The model was calibrated to the following initial yield curve:

Table 26: Initial Yield Curve for Hull-White Model Calibration

Term	Yield
0.25	2.49%
1.00	2.73%
5.00	4.03%
20.00	4.95%

The volatility parameters for the Hull-White model were set based upon historical interest rate volatility.

10.4. Calculation of the Best Estimate Liability

To calculate the best estimate liability for interest-sensitive cash flows such as these, it is common practice to draw upon methods used to price interest rate derivative securities.

Using this method, a risk-neutral interest rate model is used to stochastically generate interest rates for future periods and to estimate the cash flows conditional on these (paths) of interest rates. For each possible path, the present value is calculated using the interest rates along the path[24].

For this product, the liability cash flows depend on the crediting rates (including the impact of the crediting rate guarantee) and the impact of lapsation. To determine the BEL, Monte Carlo simulation is used to derive a sample of risk-neutral interest rate paths and then the crediting rates and lapse rates for that set of risk-free rates are estimated to determine the cash flows at each period. Discounting these cash flows at the compounded rates along each path provides a sample of discounted values. The average of this sample represents the BEL.

[24] In some instances, an analytical option pricing formula could be used in place of Monte Carlo simulation methods. But for many products with similar characteristics to this simplified SPDA product, such analytical formulas may be far less practical. Therefore, in the discussion that follows we will assume that Monte Carlo methods are used to value the embedded options.

For instance, for the product described here we established a particular set of assumptions regarding the relationship between the one-period risk free rates at each point in time, the crediting rates and the lapse rates, as described above. Using those assumptions, a simulation of 5,000 interest rate paths and liability cash flows resulted in a BEL of $100.61 million.

Our sample model allows for "dynamic lapse" depicted as a fixed functional relationship between the spread between the market rate and the guaranteed crediting rate. It also allows for two sources of stochastic variability of the actual lapse rates – a stochastic parameter within the dynamic lapse function and a stochastic noise term. As a result of these three distinct factors that drive the actual lapse rate, reliable estimates of the BEL may require an extensive number of simulation iterations *for each interest rate path*. In addition, a large number of interest rate paths, or alternatively, very careful selection of representative paths, would be needed to ensure that scenarios where the interest rate guarantee enhances the policyholder benefits are adequately reflected in the average. Even when advanced simulation methods are used, large measurement errors can occur if an insufficient number of unique paths and simulated lapse rates are used. While we believe this potential for measurement error is significant, in what follows we will assume that this can be managed effectively and that practical guidelines can be established.

With respect the interpretation of the BEL, for products such as the SPDA where the cash flows are sensitive to interest rates (or other market variables) it is not possible to cleanly separate the BEL and the MVM. As a result, the BEL shown above is not literally a pure best estimate liability. Because risk-neutral interest rate scenarios are used to project the cash flows and perform the discounting, the estimate derived *implicitly* incorporates a risk margin for the interest rate risk. While this implicit risk margin does not account for the random variation associated with the lapse rate function (e.g. parameter uncertainty, random noise term), it does incorporate the risk margin for the pure interest rate risk as well as the non-stochastic portion of the lapse function for a given interest rate scenario (e.g. the "dynamic lapse"). The additional risk margin associated with the random variation in lapse rates is what is explicitly calculated using the cost of capital method.

10.5. Calculation of the Capital Base for Non-Hedgeable Risks

For the purposes of this appendix, we will not show the complete calculation of the MVM using the cost of capital method or repeat any of the discussion of the two alternative approaches that could be used to determine the capital base using either an ultimate risk exposure horizon or a sequentially, one-year risk exposure horizon. Nor will we repeat the discussion from Section 6 regarding the adjustments needed to similarly calculate the MVM for solvency applications.

Instead, our purpose here is simply to demonstrate the challenges that arise when the capital base for non-hedgeable risks is calculated for products that also contain significant hedgeable risks. We therefore limit the following discussion to the case where an ultimate risk exposure horizon is used, though the points raised would also be applicable using a one-year risk exposure horizon.

10.5.1. Full Stochastic Simulation and the Double-Counting of Risk

To calculate the initial capital base using an ultimate risk exposure horizon, the goal is to quantify the additional capital that would have to be available immediately in the event the lapse rate variability in all future periods causes materially higher policyholder benefit payments than is expected.

One approach to quantifying this amount is to use the simulated interest rate paths and their associated liability cash flows (5,000 unique values discounted values) and then identify an extreme value such as the 99.5th percentile. The difference between this value and the BEL would appear to be comparable to the capital base previously calculated in other examples. However, there is an important caveat that limits the usefulness of this measure[25]. The BEL for a product such as this already includes a risk margin for the interest rate risk. As a result, a portion of the difference between the 99.5th percentile and the BEL simply represents the impact of interest rates independent of the lapse variability and uncertainty. Including this component would incorporate a risk margin for this risk twice.

10.5.2. Using the Forward Rate Path

One way to eliminate the potential double counting of the risk margin for interest rate risk is to identify a particular base case, deterministic interest rate path and then simply simulate the stochastic portion of the lapse rates to identify a large sample of present values conditional on the selected path. For instance, the selected path could be the current forward rates. In our example, this would produce an initial capital base of $1.35 million.

10.5.3. Adverse Rate Path

While the above estimate avoids duplicating the risk margin for interest rate risk, it may substantially understate the impact of lapse rate variability because it uses one particular interest rate path. The path selected may or may not be a path that causes the interest rate guarantee to have a significant impact on the liability flows, in which case it would not be surprising for the impact of substantially lower lapsation rates to be minimal.

For example, if instead of using the forward rates as the baseline path we used the specific path that resulted in one of the largest present value liability and then simply allowed the lapse rates to vary randomly, the estimated capital would have been $8.48 million. This is significantly higher than the result using the current forward rates.

Based on this discussion, we believe that guidance may be needed to ensure that the calculations are consistently applied. There does not appear to be any clear methodology that is entirely consistent with the intent of determining a risk margin for the non-hedgeable portion of the risk using the cost of capital methodology.

10.5.4. Lapse Rate Shock Scenarios

The above calculations used two alternative deterministic interest rate paths along with a stochastic simulation reflecting the lapse variability we noted earlier. However, it is more practical to perform these calculations using selected “shocks” to key lapse rate assumptions rather than to simulate the impacts of lapse rate uncertainty and variability. The challenge though is that this leaves one uncertain as to the risk exposure horizon and confidence level being achieved.

For example, under Solvency II’s quantitative impact studies it was proposed that all lapse rates could be shocked up or down 50%, whichever resulted in the worst outcome, subject to a minimum shock of 3% each period[26]. Applying this scenario approach instead of incorporating the stochastic variability as in our previous calculations, but still using the adverse interest rate path, the resulting capital base would have been $4.58 million, which is 46% less than the estimate shown above using the stochastic analysis.

---

[25] Another caveat, albeit more subtle, relates to the use of the so-called *risk neutral valuation methodology*. The risk neutral interest rate model used to determine the BEL properly calculates the “value” of uncertain cash flows as a simple probability weighted average of their present values along the simulated interest rate paths because of a subtle but critical adjustment to the probabilities assigned to each scenario. In addition, a variety of alternative interest rate models may produce the same value of the cash flows but may indicate substantially different probabilities for any specific scenario. As a result, it is not technically possible to identify a particular percentile, such as the 99.5th percentile, from the 5,000 simulated scenarios.

[26] For the purposes here we will apply these shocks in the same manner as is suggested in the Solvency II QIS 2 and QIS 3 Technical Documents. That is, we incorporate shock in each period over the lifetime of the liability and not simply in the first period. This approach inherently assumes that the “shocks” in all periods are perfectly correlated and appears to be inconsistent with the intention of determining the impact on the BEL over just *one year*.

10.6. Comparison of MVM Calculations

Clearly the differences in the initial capital base under each of these approaches will translate, without any other adjustments, to similar differences in the MVM calculations. Using for simplicity the assumption that the capital base each period remains the same proportion of the BEL and an annual cost of capital rate of 6%, the following alternative MVM estimates would be obtained:

Table 27: Comparison of Alternative Estimates of MVM

Method	MVM (Millions)
Stochastic Analysis Using Adverse Rate Path	2.74
Shock Scenario	1.48
Stochastic Analysis Using Forward Rate Path	0.44

It is difficult to draw definitive conclusions from the relationship between these estimates. The magnitude of the difference depends on a variety of factors, including the risk characteristics of the product, the assumptions used regarding the underlying variability and parameter uncertainty, the form of the dynamic lapse function, and the assumed independence of the lapse variability in the stochastic case compared to the assumed perfect correlation in the scenario analysis (shock) case. Nonetheless, this example highlights the need to more carefully demonstrate that the shock scenarios used are consistent with the confidence level being sought for the initial capital.

In each of these calculations, the focus was on the MVM for financial reporting purposes. Similar adjustments as those described in Section 6 would have to be made if the calculations were being performed for solvency applications.

11. Appendix D: P&C General Liability Risk

In this appendix we use a P&C general liability product example to demonstrate the implementation issues addressed in earlier sections.

11.1. Assumptions

We assume a new portfolio of general liability policies totaling $100 million of premium is written. The loss ratio, which is assumed to be lognormally distributed, has an expected value of 85% and a standard deviation of 15%. We further assume that the claims will be paid according to the following payment pattern:

Table 28: Assumed Claim Payment Pattern

Year	Claim Payment Pattern
1	5.0%
2	12.0%
3	22.0%
4	27.0%
5	12.0%
6	8.0%
7	6.0%
8	4.5%
9	3.0%
10	0.5%

These assumptions are consistent with historical industry experience from 1996 – 2004. These assumptions are likely to vary considerably in practice for companies writing similar policies when factors such as the size of the portfolio (and hence the amount of internal risk diversification across policyholders), the attachment point and policy limits profiles and other factors unique to each portfolio are taken into consideration. Attempts to mandate specific assumptions that should be used would not be appropriate and identifying instances where differences in assumptions across insurers are inconsistent with differences in their portfolios may be impossible. To ensure consistency, some standardization of model assumptions may therefore be needed[27].

We do not incorporate uncertainty into the timing of the claim payments in this example, though this would normally need to be reflected in practice. Other relevant assumptions are that the yield curve is currently flat at 4.00% per annum for all maturities and interest rates will not change. In practice the assumption of a constant yield curve would have to be modified. That modification, along with the payment pattern uncertainty, would introduce additional liquidity risk that is not hedgeable and therefore would be included within the MVM.

11.2. Best Estimate Liability

Given the assumptions specified above, the BEL can be calculated as $72.3 million. This is merely the present value of the expected claim payments using the current risk-free yield curve to determine the discount factors.

---

[27] The IAA's *Global Framework for Insurer Solvency Assessment,* Appendix B, contains one model that could be used to reflect differences in portfolio size, attachment points and policy limits and still incorporate standardized assumptions across insurers.

11.3. Market Value Margin for Financial Reporting

In Section 3 we discussed that there are two alternative perspectives that can be used to determine the capital base – one that assumes the ultimate risk must be reflected in the capital base from inception and one that assumes that capital can be funded sequentially as a series of one-year commitments. The choice among these methods can significantly impact the resulting MVM estimates unless appropriate adjustments are made to the cost of capital rate, as we will demonstrate here.

11.3.1. Ultimate Risk Exposure Horizon

In the base case, we define the capital base as the difference between the discounted 99.5th percentile loss and the BEL. Given our assumptions, the 99.5th percentile value is $111.7 and the capital base is therefore equal to $39.5 million. We also assume that the capital base in subsequent periods can be approximated by assuming that the initial capital is "released" in proportion to the claim payment pattern, as follows:

Table 29: Assumed Pattern of Capital Release – Base Case

Year	Claim Payment Pattern	Percent Unpaid Beginning of Year	Capital Released
1	5.0%	100.0%	5.0%
2	12.0%	95.0%	12.0%
3	22.0%	83.0%	22.0%
4	27.0%	61.0%	27.0%
5	12.0%	34.0%	12.0%
6	8.0%	22.0%	8.0%
7	6.0%	14.0%	6.0%
8	4.5%	8.0%	4.5%
9	3.0%	3.5%	3.0%
10	0.5%	0.5%	0.5%

Without necessarily advocating any particular annual cost of capital rate to be appropriate, we use a 6% annual rate in the example calculation below.

Combining these three sets of assumptions, the MVM is calculated as follows:

Table 30: Calculation of MVM – Base Case

Year	Capital Base Beginning of Year	Annual Rate	Nominal Cost of Capital	PV Factor	Present Value Cost of Capital
1	39,465,418	6.00%	2,367,925	0.96	2,276,851
2	37,492,148	6.00%	2,249,529	0.92	2,079,816
3	32,756,297	6.00%	1,965,378	0.89	1,747,214
4	24,073,905	6.00%	1,444,434	0.85	1,234,709
5	13,418,242	6.00%	805,095	0.82	661,729
6	8,682,392	6.00%	520,944	0.79	411,709
7	5,525,159	6.00%	331,510	0.76	251,920
8	3,157,233	6.00%	189,434	0.73	138,418
9	1,381,290	6.00%	82,877	0.70	58,229
10	197,327	6.00%	11,840	0.68	7,998
				Total MVM:	8,868,592
				% of BEL:	12.3%

This base case MVM estimate of $8.87 million assumes that the ultimate risk has to be funded at inception and uses the 99.5th percentile to define the capital base in the MVM calculation.

As we have previously discussed, some practitioners suggest that the percentile used should be adjusted to reflect the use of an ultimate risk exposure horizon. In this case, the average term of the liability is 4.2 years. Using the alternative percentiles from Fitch shown in Section 3.3.4 (e.g. 98.2%) that correspond to a 4.2 year term, the MVM would be reduced to $6.82 million, or 9.4% of the BEL.

These calculations also assume that the portfolio is assessed on a *stand-alone* basis without any additional diversification adjustments with other lines of business. Diversification of risk within the portfolio has been reflected in the model parameters, but diversification across other risks in the insurer's portfolio has not been reflected. In practice, the level of granularity would have to be specified. In some applications it may be more appropriate to include, for instance, all risks for new policies of a certain class of business with the risks associated with the loss reserves for that same class of business. In this case, additional diversification adjustments beyond those reflected here would have to be considered.

11.3.2. One-Year Risk Exposure Horizon

The base case calculations above reflected a loss ratio standard deviation of 15%, which is a reasonable amount of variation when the ultimate risk exposure horizon is used. However, over a shorter risk exposure horizon such as one year, industry data indicates substantially lower loss ratio variation. For many longer-tailed P&C risks, the ultimate liability estimates in the earliest years since policy inception are usually set at or near the expected value. It is only in later years, as actual claim experience emerges, that revisions generally occur. This can be demonstrated using the following P&C industry loss, expense and premium data taken from AM Best's Aggregates and Averages (2005) for the general liability line of business:

Table 31: Comparison of Loss Ratio Variability – General Liability (Occurrence)

		Loss and Loss Adjustment Expenses		Loss and Loss Adjustment Expense Ratio	
Accident Year	Net Earned Premiums	Initial Estimate	Current Estimate	Initial Estimate	Current Estimate
1996	11,575,802	9,272,671	9,150,431	80.1%	79.0%
1997	12,287,031	9,967,825	10,613,179	81.1%	86.4%
1998	12,937,150	10,687,120	12,699,218	82.6%	98.2%
1999	11,950,109	9,498,034	11,943,808	79.5%	99.9%
2000	12,020,345	9,564,004	11,781,294	79.6%	98.0%
2001	12,817,620	11,443,208	12,584,226	89.3%	98.2%
2002	17,173,523	12,328,318	13,135,327	71.8%	76.5%
2003	21,339,942	14,877,656	14,127,624	69.7%	66.2%
2004	24,990,169	17,019,781	15,066,190	68.1%	60.3%
			Standard Deviation:	6.8%	15.1%

Given that this data indicates loss ratio variability over the course of the first year that is substantially lower than the variability over a longer risk exposure horizon, it would not be appropriate to base the MVM calculation on this lower estimate of variability without also incorporating further adjustments to the method for estimating the risk measure in future periods, the annual cost of capital rate, or both. Without these adjustments, the resulting MVM would decline considerably to 5.0% of the BEL (compared to 12.3% of the BEL in the base case and 9.4% using adjusted percentile levels), as shown below:

Table 32: Calculation of MVM – One-Year Risk Exposure Horizon

Year	Capital Base Beginning of Year	Annual Rate	Nominal Cost of Capital	PV Factor	Present Value Cost of Capital
1	16,225,071	6.00%	973,504	0.96	936,062
2	15,413,817	6.00%	924,829	0.92	855,056
3	13,466,809	6.00%	808,009	0.89	718,317
4	9,897,293	6.00%	593,838	0.85	507,615
5	5,516,524	6.00%	330,991	0.82	272,051
6	3,569,516	6.00%	214,171	0.79	169,262
7	2,271,510	6.00%	136,291	0.76	103,570
8	1,298,006	6.00%	77,880	0.73	56,906
9	567,877	6.00%	34,073	0.70	23,939
10	81,125	6.00%	4,868	0.68	3,288
				Total MVM:	3,646,066
				% of BEL:	5.0%

The differences between these two calculations can be reduced to some degree through the use of more sophisticated assumptions regarding the relationship between the initial capital base and the capital base in subsequent periods. However, it would not be possible to obtain the *same* estimate of the MVM as in the base case unless the annual cost of capital rate is modified as well.

11.4. Market Value Margin for Solvency Applications

For solvency applications, the calculation of the MVM follows essentially the same procedure. All of the considerations discussed above apply as well in this case. However, as described in Section 6, there are additional considerations that cause important differences in the implementation of the cost of capital methodology.

One important difference is that the first period's contribution to the MVM would not be included. For instance, if the ultimate risk exposure horizon is used and the first period's contribution to cost of capital is ignored, the following would be obtained for the distress scenario MVM:

Table 33: Calculation of MVM – Solvency Application

Year	Capital Base Beginning of Year	Annual Rate	Nominal Cost of Capital	PV Factor	Present Value Cost of Capital
1	N/A	6.00%	0	0.96	0
2	37,492,148	6.00%	2,249,529	0.92	2,079,816
3	32,756,297	6.00%	1,965,378	0.89	1,747,214
4	24,073,905	6.00%	1,444,434	0.85	1,234,709
5	13,418,242	6.00%	805,095	0.82	661,729
6	8,682,392	6.00%	520,944	0.79	411,709
7	5,525,159	6.00%	331,510	0.76	251,920
8	3,157,233	6.00%	189,434	0.73	138,418
9	1,381,290	6.00%	82,877	0.70	58,229
10	197,327	6.00%	11,840	0.68	7,998
				Total MVM:	6,591,741

Notice that in this calculation, as in the fair value financial reporting case, the ultimate risk exposure horizon is used to calculate the distress scenario MVM. This is done even though the change in the BEL (referred to as the SCR) is measured over a one year solvency time horizon.

Under Solvency II and the Swiss Solvency Test, a one-year risk exposure horizon is advocated for the purposes of calculating the MVM. In those implementations of the cost of capital methodology, the case is made that using a one-year risk exposure horizon for the MVM calculations simplifies the implementation of the cost of capital method "due to the fact that there is only one unknown, the SCR for non-hedgeable risks, and this can be calculated with ease using the standard SCR [formulas]."[28]

We acknowledge the practical appeal of following this approach but nonetheless refer to our earlier concerns about the proper calculation of the capital base in future periods and the proper calibration of the annual cost of capital rate when the one-year risk exposure horizon is used. As discussed in Section 6.2, other adjustments to the MVM calculation may also be needed to take into account the conditions that might apply in a distress scenario, such as changes to the cost of capital rate or changes in the perceived risk.

[28] Chief Risk Officer Forum, *A Market Cost of Capital Approach to Market Value Margins – Discussion Paper*, March 17, 2006.

12. Appendix E – Approximating the Capital Base Over Time

12.1. Introduction

This appendix demonstrates the potential inaccuracy of approximating the capital base in future periods for the MVM calculation using a one-year risk exposure horizon and applying a constant ratio to the annual outstanding reserve balances.

12.2. Data

The example uses US commercial auto liability industry data from Aggregate Industry Schedule P. A variety of methods exist for both estimating loss reserves and for quantifying the uncertainty around those estimates. For the purpose of this example, we use the methods developed by Ben Zehnwirth[29], which performs both of these tasks simultaneously using a robust probabilistic model of the underlying claim process.

Applying those methods to the industry loss reserve data produces the following table of incremental claim payments, by accident year, for several subsequent calendar periods.

Table 34: Estimated Reserves and their Variability – Commercial Auto Liability

Accident Period vs Development Period																			
	Cal. Per. Total	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Reserve	Ultimate
1997	2,308,574	2,308,574	2,548,125	2,153,308	1,537,792	894,002	453,769	191,576	83,961	40,476	17,768	7,806	3,432	1,510	665	293	129	13,835	10,146,160
	2,354,733	2,354,733	2,522,414	2,109,939	1,458,759	855,372	455,584	197,939	87,631	81,371	8,583	3,831	1,720	775	351	159	72	4,633	4,633
1998	4,899,477	2,351,352	2,595,400	2,193,310	1,566,402	874,324	443,820	195,130	85,520	41,229	18,099	7,951	3,496	1,538	677	299	132	32,192	10,377,621
	4,878,597	2,356,183	2,624,303	2,137,469	1,554,034	902,836	447,191	196,154	101,450	25,809	8,734	3,904	1,753	790	357	162	74	10,341	10,341
1999	7,287,884	2,539,176	2,802,769	2,368,597	1,624,158	906,650	479,274	210,720	92,354	44,524	19,545	8,587	3,775	1,661	732	322	142	79,289	10,999,828
	7,237,104	2,502,862	2,857,540	2,355,317	1,657,291	903,049	334,037	239,915	70,528	21,202	9,435	4,218	1,894	854	386	175	80	24,604	24,604
2000	9,120,260	2,586,388	2,854,946	2,316,529	1,588,617	923,471	488,182	214,642	94,075	45,355	19,911	8,748	3,846	1,693	745	329	145	174,846	11,244,901
	9,104,360	2,650,592	2,913,575	2,321,575	1,561,428	909,624	510,292	202,969	10,325	21,607	9,616	4,300	1,931	870	394	179	81	28,556	28,556
2001	10,168,898	2,484,951	2,633,600	2,137,141	1,526,191	887,218	469,029	206,228	90,391	43,581	19,132	8,406	3,696	1,627	716	316	139	374,233	10,813,045
	10,279,111	2,600,813	2,770,871	2,230,771	1,528,984	870,567	436,806	20,748	9,969	20,772	9,247	4,135	1,857	837	379	172	78	36,416	36,416
2002	10,332,651	2,430,271	2,575,935	2,176,794	1,554,550	903,731	477,774	210,079	92,081	44,397	19,491	8,564	3,766	1,657	730	322	142	859,003	10,439,134
	10,394,282	2,286,125	2,475,640	2,284,174	1,588,708	945,484	20,579	21,310	10,246	21,176	9,428	4,217	1,894	854	387	176	80	43,919	43,919
2003	10,220,783	2,377,044	2,623,746	2,217,250	1,583,483	920,579	486,698	214,008	93,806	45,229	19,857	8,725	3,837	1,688	744	328	145	1,795,644	10,454,998
	10,087,587	2,271,569	2,564,154	2,194,280	1,629,351	38,321	21,936	21,948	10,556	21,592	9,615	4,301	1,932	871	394	179	82	64,199	64,199
2004	10,429,749	2,421,183	2,672,526	2,258,526	1,613,004	937,770	495,803	218,018	95,566	46,079	20,231	8,889	3,909	1,720	758	334	147	3,442,228	10,561,893
	10,329,024	2,424,266	2,519,661	2,175,738	66,524	40,896	23,612	22,665	10,898	22,021	9,807	4,388	1,971	889	402	183	83	106,172	106,172
2005	10,622,657	2,466,216	2,722,294	2,300,640	1,643,124	955,311	505,094	222,109	97,362	46,946	20,612	9,057	3,983	1,753	772	340	150	5,807,252	11,028,226
	10,541,181	2,434,937	2,786,037	94,822	71,049	44,105	25,583	23,462	11,273	22,462	10,005	4,477	2,012	907	411	187	85	170,273	170,273
2006	10,815,218	2,512,162	2,773,073	2,343,610	1,673,858	973,209	514,574	226,283	99,194	47,831	21,001	9,228	4,058	1,786	787	347	153	8,688,992	11,284,671
	10,876,984	2,595,679	114,913	101,478	76,689	47,893	27,822	24,341	11,683	22,917	10,209	4,568	2,053	926	419	190	87	254,539	254,539
	Total Fitted/Paid		2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	Total Reserve	Total Ultimate
Cal. Per.	86,206,151		8,455,802	5,787,956	3,508,146	1,868,382	911,853	404,693	181,687	84,013	36,846	16,136	7,038	3,037	1,274	497	153	21,267,514	107,350,477
Total	86,082,963		202,665	165,857	120,414	79,197	53,478	41,623	30,437	26,591	12,086	5,522	2,530	1,156	521	226	87	519,282	519,282

The blue values in the upper left portion of the table are the actual incremental payments and the black values are the fitted values based on the reserve model. The black values in the "forecast" portion of the table (lower right) are the expected claim payments by accident year and by calendar year. Along any given diagonal the values represent the calendar year payments. The total for the calendar period is then summarized on the row second from the bottom.

The red values are the standard deviations of each cell and the aggregate standard deviations, either across the rows for an accident year or across a diagonal for a calendar year, and reflect the model's estimated correlation structure (i.e. they are not simple sums of the cell elements).

12.3. Estimated Reserves and their Variability

What this example shows is that the aggregate estimated *ultimate* loss for these accident years is 107,350,477 and the estimated reserves are 21,267,514. The standard deviation, which is the same in dollars whether discussing the reserve or the ultimate because the paid to date is not random, is 519,282. This gives a coefficient of variation (the ratio of standard deviation to the mean) for the reserves equal to 2.44%. This is a measure of the *ultimate* variability for the entire portfolio of reserves at a point in time.

[29] See Barnett, G.; and Zehnwirth, B, "Best Estimates for Reserves," *PCAS* LXXXVII, 2000, pp. 245-303.

An alternative metric is the sequence of calendar year payments – their expected values and their standard deviations. These are shown, aggregated across all accident years, in the bottom rows and summarized in the table below:

Table 35: Standard Deviation of Annual Payments

Calendar Year	Mean Payment	Standard Dev.	CV
2007	8,455,802	202,665	0.020
2008	5,787,956	165,857	0.030
2009	3,508,146	120,414	0.030
2010	1,868,382	79,197	0.040
2011	911,853	53,478	0.060
2012	404,693	41,623	0.100
2013	181,687	30,437	0.170
2014	84,013	26,591	0.320
2015	36,846	12,086	0.330
2016	16,136	5,522	0.340
2017	7,038	2,530	0.360
2018	3,037	1,156	0.380
2019	1,274	521	0.410
2020	497	226	0.460
2021	153	87	0.570

A variation on this table that is more useful for the present discussion is to show the variability relative to the beginning of year loss reserves rather than the calendar year payments. This is shown as follows:

Table 36: Standard Deviation of Annual Payments

(Relative to Outstanding Reserves)

Calendar Year	BOY Reserve	Standard Dev.	Reserve CV
2007	21,267,513	202,665	0.010
2008	12,811,711	165,857	0.013
2009	7,023,755	120,414	0.017
2010	3,515,609	79,197	0.023
2011	1,647,227	53,478	0.032
2012	735,374	41,623	0.057
2013	330,681	30,437	0.092
2014	148,994	26,591	0.178
2015	64,981	12,086	0.186
2016	28,135	5,522	0.196
2017	11,999	2,530	0.211
2018	4,961	1,156	0.233
2019	1,924	521	0.271
2020	650	226	0.348
2021	153	87	0.569

12.4. Best Estimate Loss Reserve Variability Over "One Year"

When a one-year risk exposure horizon is used it is necessary to quantify the potential change in the best estimate liability over a one year period and to determine the sequence of these one-year measures over the run-off period of the liabilities. The previous table of calendar year payment variability does not necessarily represent this quantity. These calendar year variability measures assume that variation in any given calendar period does not impact the mean value or standard deviations of the *remaining* calendar year payments. This would be true if all of the variability was assumed to represent statistical variation and that the party estimating the future payments interpreted it as such.

In practice though, adverse or favorable experience in any given period (or over several periods) may cause the analyst or a third party to re-evaluate the reserve estimation model and/or its parameters. This would introduce a secondary effect that could either dampen or worsen the estimates of the remaining payments.

As a simplifying assumption, these secondary effects can perhaps be ignored. In this case, the variability of the calendar year payments would directly translate into variability in the best estimate liability.

Further, the sequence of one-year measures of variability represent current estimates conditional on future experience emerging as *expected*. This could be different from a set of estimates of the future levels of variability that are conditional on adverse results being obtained in each subsequent period. Care would have to be taken to ensure that the sequence is obtained in a consistent fashion across different practitioners.

12.5. Relationship Between the Sequence of One-Year Variability Measures

Under the Solvency II QIS 3 process, one example of a practical implementation of the cost of capital methodology, an approximation is recommended whereby the sequence of one-year variability measures is estimated using the initial estimate and the current reserves to obtain a ratio that can be applied to the reserve balances in all subsequent periods.

These data can be used to test the appropriateness of this approximation. As the table below shows, this approximation is quite poor:

Table 37: Comparison of Approximate and Modeled Sequence of Estimated Variability

Initial Reserves	21,267,513
Initial Standard Deviation	202,665
Ratio	0.00953

Calendar Year	BOY Reserve	Approximate Std. Dev	Model Std. Dev
2007	21,267,513	202,665	202,665
2008	12,811,711	122,087	165,857
2009	7,023,755	66,932	120,414
2010	3,515,609	33,501	79,197
2011	1,647,227	15,697	53,478
2012	735,374	7,008	41,623
2013	330,681	3,151	30,437
2014	148,994	1,420	26,591
2015	64,981	619	12,086
2016	28,135	268	5,522
2017	11,999	114	2,530
2018	4,961	47	1,156
2019	1,924	18	521
2020	650	6	226
2021	153	1	87

12.6. Relationship Between the Sequence of Capital Measures

The previous table understates the impact of these estimation errors because it focuses on the standard deviation rather than the *capital base* that would result from the full risk measure calculation.

Assuming a lognormal distribution, the formula used to convert the coefficient of variation into a risk measure at the 99.5th percentile is:

$$\rho(\sigma) = \frac{\exp\left[\Phi^{-1}(.995)\sqrt{\ln(\sigma^2+1)}\right]}{\sqrt{(\sigma^2+1)}} - 1 \approx 3\sigma$$

Ignoring diversification effects across other risk types, the following comparison would be obtained using either the proposed approximation or the model standard deviations by period:

Table 38: Comparison of Approximate and Modeled Sequence of Capital Base Estimates

Calendar Year	BOY Reserve	Approximate Std. Dev	Model Std. Dev	Approximate Capital	Capital
2007	21,267,513	202,665	202,665	527,488	527,488
2008	12,811,711	122,087	165,857	317,763	433,294
2009	7,023,755	66,932	120,414	174,207	316,014
2010	3,515,609	33,501	79,197	87,196	209,060
2011	1,647,227	15,697	53,478	40,855	142,691
2012	735,374	7,008	41,623	18,239	113,965
2013	330,681	3,151	30,437	8,202	86,502
2014	148,994	1,420	26,591	3,695	82,450
2015	64,981	619	12,086	1,612	37,750
2016	28,135	268	5,522	698	17,419
2017	11,999	114	2,530	298	8,092
2018	4,961	47	1,156	123	3,775
2019	1,924	18	521	48	1,761
2020	650	6	226	16	816
2021	153	1	87	4	367
			PV Capital:	1,089,237	1,748,412
			Cost of Capital:	6.00%	6.00%
			MVM:	65,354	104,905

Once again, this shows that the approximation would understate the market value margin (MVM), producing an estimate that is only 62% of the results using the full model.

12.7. Other Lines of Business

Similar results can be obtained for other lines of business as well. In some cases, the differences in results are even more extreme. For instance, applying the same analysis against industry GL (Occurrence) data results in an approximate MVM that is only 36% of the results obtained using the full model.

These results are shown in the following summary table for this line of business.

Table 39: Comparison of Approximate and Modeled MVM – General Liability

Calendar Year	BOY Reserve	Approximate Std. Dev	Model Std. Dev	Approximate Capital	Model Capital
2007	41,068,841	420,540	420,540	1,095,412	1,095,412
2008	31,746,805	325,083	413,027	846,769	1,079,090
2009	24,074,139	246,516	376,216	642,119	985,715
2010	18,332,950	187,727	342,844	488,987	901,279
2011	14,198,985	145,396	334,588	378,724	884,221
2012	11,147,186	114,146	323,567	297,324	860,156
2013	8,830,720	90,426	335,072	235,538	899,330
2014	7,007,575	71,757	334,149	186,910	906,257
2015	5,568,926	57,025	322,660	148,538	884,708
2016	4,430,702	45,370	304,138	118,178	843,426
2017	3,527,803	36,124	281,659	94,096	790,339
2018	2,809,695	28,771	257,495	74,942	731,496
2019	2,237,057	22,907	233,190	59,668	671,148
2020	1,779,222	18,219	209,731	47,456	612,140
2021	1,412,213	14,461	187,704	37,667	556,283
2022	1,117,238	11,440	167,415	29,800	504,659
2023	879,536	9,006	148,985	23,459	457,872
2024	687,484	7,040	132,414	18,337	416,230
2025	531,907	5,447	117,627	14,187	379,906
2026	405,546	4,153	104,508	10,817	349,085
2027	302,645	3,099	87,802	8,072	301,852
2028	221,620	2,269	72,954	5,911	259,316
2029	158,321	1,621	59,760	4,223	220,896
2030	109,423	1,120	48,036	2,919	186,009
2031	72,255	740	37,609	1,927	153,975
2032	44,668	457	28,326	1,191	123,975
2033	24,925	255	20,048	665	94,877
2034	11,623	119	12,653	310	65,045
2035	3,623	37	6,033	97	32,668
			PV Capital:	4,122,188	11,298,168
			Cost of Capital:	6.00%	6.00%
			MVM:	247,331	677,890

13. Appendix F – Shareholder Expected Return Models

This appendix discusses some of the practical parameter estimation decisions that need to be made in order to apply the Capital Asset Pricing Model (CAPM) or the Fama-French 3-Factor Model to quantify risk-adjusted shareholder return expectations.

13.1. Quantification Using CAPM

The CAPM assumes that the risk premium over the risk free rate is a function of two key inputs, the measure of systematic risk (beta) and the expected equity market risk premium as measured by the expected overall equity market return relative to the risk free rate. Neither of these quantities is readily observable and both are measured with considerable uncertainty.

With respect to the beta, two issues would need to be addressed. First, should company-specific betas or industry-wide betas be used? If company-specific betas are used the degree of measurement error and period to period volatility would be substantially greater than if industry-wide beta estimates were used. Second, what specific methods, time periods and statistical adjustments should be made and how often should the estimates be updated?

With respect to the equity market risk premium, similar estimation error exists, though its impact could be even greater. For example, a recently completed exhaustive survey of numerous academic and practitioner-oriented estimates of the equity market risk premium[30] produced estimates that ranged from as low as 2% to as high as 8% per annum. Given the long term nature of the insurance liabilities and the need to earn the target spread over the lifetime of the capital commitment, this range of estimates would lead to significant differences in the reported risk margins.

Finally, we note that there are strong arguments suggesting that CAPM beta estimates should vary by product, though quantification of product-specific betas introduces significant additional challenges and introduce yet another level of uncertainty.

To appreciate the range of MVM estimates that could be obtained, the following table shows the results for two different estimates of an industry beta and three different estimates of the equity market risk premium.

- The two beta estimates reflect historical betas for the P&C insurance industry from two different sources. One is based on research performed by Cummins and Phillips[31] using 60 months of data to obtain quarterly beta estimates from 1997 – 2000 and incorporating a "sum-beta" adjustment. The other is based on research performed by Swiss Re[32] which uses a long-term average of two-year beta estimates from 1977 to 2003 and includes an adjustment to reflect distortions in the S&P 500 index caused by the technology bubble from 1999-2002.
- The three equity market risk premiums reflect three different estimates referenced in the Derrig and Orr study. One reflects a conditional estimate that takes into account current market conditions, one reflects the same underlying assumptions as the conditional estimate but has been adjusted to reflect an unconditional or long-range average estimate and the third merely reflects a historical average.

[30] See Derrig & Orr

[31] See Cummins & Phillips.

[32] See Swiss Re Sigma.

The range of spreads to the risk free rate that these assumptions produce is shown below:

Table 40: Range of Cost of Capital Spreads to Risk Free Rate - P&C Insurance Industry Average

		Equity Risk Premium		
		Conditional	Unconditional	Historical
Source	Industry Beta	***2.40%***	***5.00%***	***8.40%***
Cummins & Phillips	***0.836***	2.01%	4.18%	7.02%
Swiss Re Study	***1.030***	2.47%	5.15%	8.65%

13.2. Quantification Using Alternative Models

The CAPM adopts a rather limited definition of systematic risk as being solely reflected in the sensitivity of returns to overall market returns (the beta). Academic and practitioner dissatisfaction with this model has led to considerable support for variations of the CAPM model that incorporate additional sources of systematic risk, such as the Fama-French 3-Factor Model that incorporates two additional risk factors. Although we do not believe that this alternative model is in widespread use, it can result in considerably larger MVM estimates.

In a recent research paper by Cummins and Phillips, an exhaustive analysis of US-based P&C insurers was used to quantify the parameters for the Fama-French 3-Factor model. The following table demonstrates the calculation of the spread to the risk free rate that results from using the Cummins and Phillips results for the overall industry. For consistency with the CAPM results shown above, we assume the same historical equity market risk premium of 8.4% and the same risk free rate of 4.0%:

Table 41: P&C Industry Average Cost of Capital Spread to Risk-Free Rate – Fama-French Model

	Industry
Risk Free Rate	4.00%
Equity Risk Premium	8.40%
Size Risk Premium	2.35%
Financial Distress Risk Premium	3.85%
Market Value (Millions)	1,910
Book Value (Millions)	979
Market to Book Ratio	1.950
Equity Beta	1.064
Size Beta	0.304
Financial Distress Beta	1.172
Estimated Required Equity Return	18.17%
Spread to Risk-Free Rate	14.17%

Notice that this estimate is considerably larger than the highest CAPM estimate because of the inclusion of two additional systematic risk factors.

13.3. Implications for Estimates of the MVM

These alternative estimates of the rate of return *per annum* will produce considerable differences in the MVM's, particularly for insurance risks that require longer term capital commitments.

Consider first the range of estimates for the expected shareholder return using the CAPM. In the P&C liability example shown in Appendix B, we used an arbitrarily 10% total cost of capital (equivalent to a 6% spread) to calculate the MVM of $7.67 million. The following table shows the range of MVM estimates that would be produced using each of the possible values for the cost of capital spread from Table 40 above:

Table 42: Range of MVM for P&C Liability Example – Alternative Beta and ERP Estimates

	Equity Risk Premium		
	Conditional	Unconditional	Historical
Beta	***2.40%***	***5.00%***	***8.40%***
0.836	2.57	5.34	8.98
1.030	3.16	6.58	11.06

This shows that depending upon the CAPM parameters used for the beta and the equity risk premium, the MVM for this particular product could vary from approximately $2.57 million to $11.06 million. If instead the spread determined using the Fama-French model was used, the MVM would rise to $17.30 million.

14. Appendix G – Tax Adjustment for MVM Calculation

The following demonstrates the process that can be used to gross up the MVM calculation under the cost of capital method to reflect the impact of corporate taxes on the risk margin.

Assume that C_i reflects the capital base each period, *h* reflects the total return demanded by the providers of this capital each period, *i* is the annual rate of investment income on the capital and d_i reflects the present value factor applied to the cost of capital each period.

In the absence of corporate taxes on the MVM and the investment income earned on the capital, the MVM calculation is as shown below:

$$MVM = \sum C_i [h - i] d_i$$

The after-tax MVM, assuming that the MVM is indeed treated as taxable income at inception, is equal to the after-tax net spread over the lifetime of the capital commitment, as shown below:

$$MVM(1-t) = C_1[h - i(1-t)]d_1 + C_2[h - i(1-t)]d_2 + \cdots + C_N[h - i(1-t)]d_n$$

$$= [h - i(1-t)] \sum C_i d_i$$

$$MVM = [\frac{h}{1-t} - i] \sum C_i d_i$$

From the final step, it is clear that to gross up the MVM to reflect the impact of corporate income taxes, the *spread* used should be $s^* = h/(1-t) - i$ instead of $s = h - i$.

15. Appendix H – Stability of Risk Margins in Various Markets

This Appendix examines the stability of risk margins in the market for S&P 500 index options and the market for credit default swaps.

15.1. S&P 500 Index Options

Short-dated put and call options on the S&P 500 index are among the most liquid options traded in the financial markets. Their prices can be used to create what is sometimes referred to as an "investor fear gauge" that indicates the magnitude and price of volatility in the index.

The following chart[33] displays the historical VIX volatility index from 1997 – 2007.

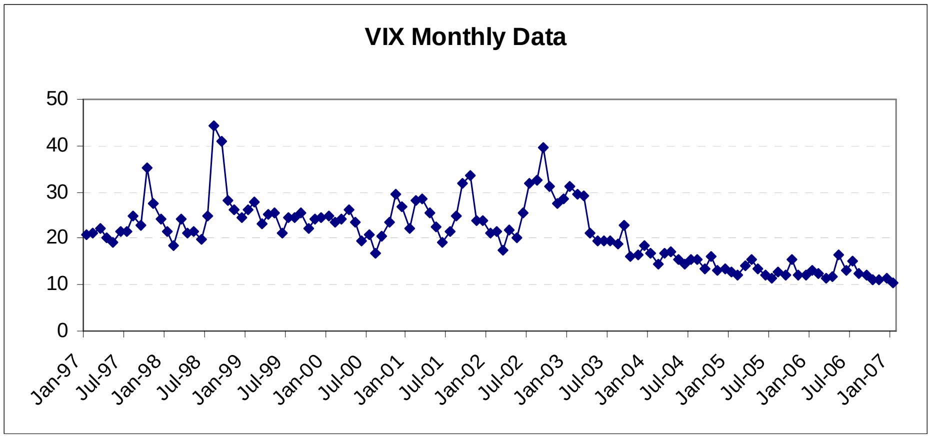

As the graph indicates, the VIX has fluctuated significantly over the period shown here, with specific periods of extreme values. This shows that rapid and substantial changes in the price of risk occur routinely even for transparent and liquid financial risks.

15.2. Credit Default Swaps

A relatively recent development in the past decade has been the enormous growth of an active and relatively liquid market for corporate bond default risk. Historically this risk was embedded in corporate bonds, but the use of credit default swaps (CDS) now allows this risk to be separated from the underlying bond or even traded outright. The pricing of CDS thus gives another perspective of how the market pricing for default risk can fluctuate over time.

As the graph below shows, the prices at which two parties have traded credit risk fluctuated significantly, dropping by nearly 50% over a relatively brief period:

[33] Source: Bloomberg. The original VIX relied upon the Black-Scholes option pricing model to estimate the implied volatility based on actual market prices for short-dated put and call options on the S&P 100 index. The new VIX is conceptually similar but does not rely on the Black-Scholes model and is based on the S&P 500 index. In either case though, the VIX is an indication of the risk margins embedded in S&P index options.

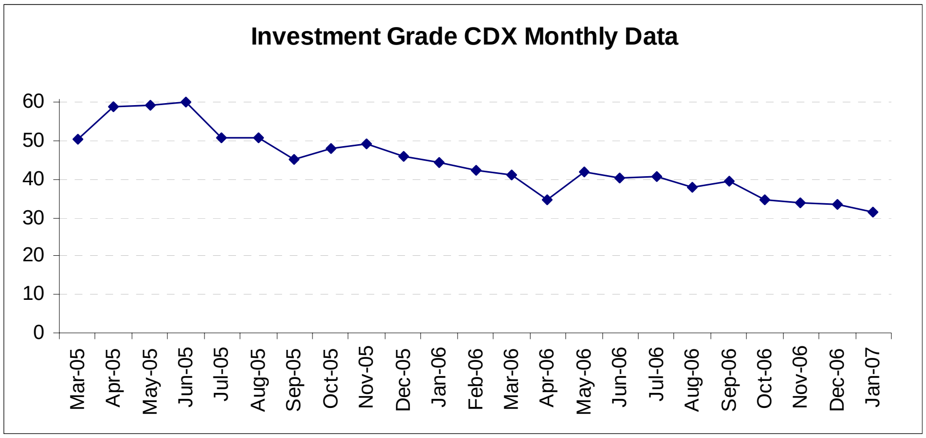

This index combines the net effect of the market's view of a) the expected default losses, b) the magnitude of the risk of substantially larger default losses than the expected value and c) the cost of the risk assumed. The changes in the index therefore reflects changes in any and call of these components, though using a diversified index suggests that the expected default losses may be relatively stable and so the variability of the costs may largely reflect changing market perceptions of the risk and/or its cost.

Given the size of the CDS market, the wealth of data available and the liquidity of the CDS market, this historical experience suggests that the stability of market value margins for insurance risk is likely to be much less than for the CDS market.

16. References

A.M. Best's, *2005 P&C Aggregates and Averages*, 2006.

Barnett, G.; and Zehnwirth, B, *Best Estimates for Reserves*, PCAS LXXXVII, 2000, pp. 245-303.

Chief Risk Officer Forum, *A Market Cost of Capital Approach to Market Value Margins – Discussion Paper*, March 17, 2006.

CIEOPS, *QIS 2 Technical Specifications*, 2006.

CIEOPS, *QIS 3 Technical Specifications*, 2007.

Cummins, J.D. and R.D. Phillips, *Estimating the Cost of Equity Capital for Property-Liability Insurers*, The Journal of Risk and Insurance, 2005, Vol. 72, No. 3, 441-478.

Derrig, Richard A and Elisha D. Orr, *Equity Risk Premium: Expectations Great and Small*, Casualty Actuarial Society Forum, 2004, Vol: Winter, Page(s): 1-44.

Federal Office of Private Insurance, *The Swiss Experience with Market Consistent Technical Provisions - the Cost of Capital Approach*, March 28, 2006.

Federal Office of Private Insurance, *SST White Paper*, 2004.

Feldblum, S., *Fair Value Accounting For Property-Casualty Insurance Liabilities*, CAS 2006 Discussion Paper Program – Current Issues in Insurance Financial Statements, 2006.

Fitch, *Exposure Draft: Prism – Insurance Rating Calibration Measures*, July 17, 2006.

Goldfarb, Richard, *Solvency II and CRO Forum Update: Market Cost of Capital Implementation*, SOA Advanced Financial Reporting Seminar, June 8-9, 2006.

Grondin, Tom, "Economic Capital Models and Implied Ratings". *Risks and Rewards*, Society of Actuaries, Issue 49, February 2007.

IAA Insurer Solvency Assessment Working Group, *A Global Framework for Insurer Solvency Assessment*, 2004.

IAIS, *The IAIS Common Structure for the Assessment of Insurer Solvency*, February 14, 2007.

Swiss Re Sigma, *Insurers' Cost of Capital and Economic Value Creation: Principles and Practical Implications*, 3/2005.