Risks and Rewards Newsletter

July 2001 – Issue No. 37

RISKS and REWARDS

The Newsletter of the Investment Section of the Society of Actuaries

NUMBER 37 JULY 2001

Exploring C1 Risk

by Thomas Merfeld

Editor's Note: This is part one of a two-part article. The second part will run in the next issue of Risks and Rewards.

Most of us consider insurance companies to be expert risk managers. One of these risks reflects the possibility that their investments perform poorly. We call this C1 risk.

I've spent years wondering how to articulate the possibility that investments perform poorly. Is an investment that you mark-to-market on the statutory filing riskier than if you could hold it at historical cost? Are private placements riskier because they don't enjoy a ready market? Are derivatives risky? Are stocks riskier in the short run than over long investment horizons? How do you isolate C1 risk from C3 risk? What is a sufficient asset reserve? Should product managers care if returns fall short of pricing assumptions? Are bond defaults worse than other causes of bond value declines? Should a P&C company own commercial mortgages? How much risk is enough? Does the character of return—income versus capital appreciation—matter? Should stocks back reserves? How bad can things get?

(continued on page 4)

CIA Task Force on Segregated Fund Investment Guarantees

excerpt from the Canadian Institute of Actuaries

Editor's Note: The CIA Task Force on Segregated Fund Investment Guarantees was founded in 1999 and charged with developing recommended approaches for the use of stochastic techniques to measure the obligations created by segregated fund investment guarantees (i.e., where an underlying level of investment performance is guaranteed by an insurer). The Task Force issued a 64-page report in August 2000 and recommended that Canadian actuaries use stochastic techniques to establish liabilities for these guarantees. The following passage on investment return models is excerpted from Section 2 of the report, and should be of particular interest to readers of this newsletter. The full report is available at the CIA Web site as accession number 20020. Also see the announcement on page 34 of this issue for the 2001 Symposium on this subject.

Policy liabilities for segregated funds, as for other policy liabilities, should be based on a prospective analysis of asset and liability cash flows. Because of the uncertainty of the underlying investment returns on which the liability costs and revenues are based, a stochastic approach is required to estimate these values.

(continued on page 9)

In This Issue

Modelling

continued from page 1

Due to the complex nature and diversity of the segregated fund guarantee, there are currently no generalized closed form solutions available to calculate the policy liabilities. A more flexible approach to calculating policy liabilities is to first use stochastic simulation to generate multiple paths of investment returns based on a selected investment return model, and then to evaluate the liability costs and revenues using the generated path set. Each path is also commonly referred to as an investment scenario or a scenario.

In the pages that follow we have separate sections that deal with the investment return models and liability cash flow models, as well as a separate discussion on the modelling of any asset hedging of these features.

2.1 Investment Return Models

A key component of stochastic modelling of the future costs associated with segregated fund investment guarantees is the model(s) used to determine investment return paths.

The task force does not support mandating specific models for establishing investment return paths. We can find no precedent for mandating specific models and we believe that such an approach would risk failure because of resistance from the membership. Instead, the task force believes that a framework which requires mandatory calibration of equity based models to specified criteria plus specific guidance/prescription that addresses certain model building items (including assumptions) can acceptably narrow the range of practice and ensure appropriate policy liabilities.

Specifically:

- Guidance is given to narrow the range of practice on setting investment assumptions.

- Investment scenario models used for the generation of equity returns will need to produce investment path results that calibrate to certain statistical criteria that measure items such as dispersion of paths and thickness of the distribution's tail.

2.1.1 Key Considerations in Selection/Development of Investment Return Models

There are a large number of investment return models and no single model can currently be identified as superior to all others. Due to the large amount of research currently going on in actuarial science, finance, econometrics, statistics and mathematics, stochastic modelling is constantly evolving. Also, due to the increasing power of computers, models that were once considered too complex to be practical can now be used. This evolution will surely continue in the future.

Notwithstanding this diversity of models, there are some requirements that need to be met in the context of using stochastic models to calculate the policy liabilities and minimum capital.

a) *Random Number Generator*

The random numbers generated by computer algorithms are called pseudo-random because they are not truly random. Knowing the algorithm and the seed to the sequence is usually sufficient to predict the next random number that will be generated.

Before using a pseudo-random number generator for stochastic simulation it should be confirmed that the generator does not exhibit any bias. This can be verified by statistical testing. The "periodicity" of the generator is the number of values that can be produced before the sequence repeats or begins to exhibit obvious bias. Some commercial software applications include pseudo-random number generators with a very low periodicity for certain seeds.

Results from stochastic modelling should be reproducible. This would ordinarily be accomplished by priming the random number generator with a "seed" value.

Variance reduction techniques can be used provided it can be demonstrated that they do not introduce any bias. It should be noted that most variance reduction techniques are designed to improve efficiency of an estimate of the mean. Where the objective is a measure of the risk arising from one tail of a distribution, some methods may in fact reduce efficiency relative to straight simulation.

b) *Number of Scenarios*

To offer some guidance as to the number of scenarios that need to be generated, recall that the standard error of the result can be expressed as a function of the square root of the number of observations. To increase the precision of the policy liability calculation, it may be necessary to increase the number of scenarios quite significantly.

The number of scenarios should be at least 1,000. The exact number to use will depend on how the scenarios will be used (e.g., calculating percentiles will generally require more scenarios than calculating expected values), and the materiality of the results. The actuary should test that the number of scenarios used provides an acceptable level of precision.

c) *Frequency*

Use of an annual projection frequency is generally acceptable for benefits/features that are not sensitive to projection frequency. The lack of sensitivity to projection frequency should be validated by testing.

Use of a more frequent projection such as a monthly frequency should always be used when product features that are sensitive to projection period frequency are present (e.g., many older age death benefits, most re-set features, etc.).

It is important that the projection frequency of the investment return model be linked appropriately to the projection period in the liability model.

Care must be taken in simulating the fee income as a percentage of the

(continued on page 10)

Modelling

continued from page 9

segregated fund. A difference in the frequency of charging the fee income and the frequency of projection of the fund could lead to an over-appraisal of the fee income.

d) *Model*

Investment returns would normally be generated on a gross basis: before the application of any fees or consideration of specific product features. The objective is to model the asset returns independently of any product features. However, care must also be taken to assess if total returns (including reinvestment of income) or price returns are required for the specific segregated funds that will be modelled.

There are a large number of potential models available and we do not want to restrict the use of any model that reasonably fits the historical data. The calibration criteria are defined below.

The model should be based on a P-measure (real world experience based valuation) as opposed to a Q-measure (risk-neutral capital markets valuation).

The P-measure approach produces a distribution of outcomes based on a real world view of outcomes for the actual assets/liabilities on the balance sheet. It generally uses historic returns and volatilities for the asset class(es) being modelled to generate investment paths. It is, therefore, consistent with the overall Canadian approach to valuation as this values liabilities in the context of the cash flow outcome on the assets currently being held and anticipated to be held in the future.

The Q-measure is appropriate in the context of financial market pricing but can produce an inappropriate valuation if the intention is not to hedge the risk using capital markets instruments. This is because it values the risk using an external capital markets framework that is independent of the expected outcomes of the actual balance sheet values being held. The Q measure approach is based on a risk neutral return framework and current investment market implied volatilities. These parameters therefore embed a significant market risk premium for absorbing the risk, particularly where there is a thin market in hedging vehicles (e.g., many long duration hedges). In addition, there is generally a lack of appropriate hedging vehicles that efficiently match the risk of many of the common design features embedded in segregated fund guarantees. This makes the derivation of an appropriate market based pricing basis difficult.

As Canadian actuarial practice implies, policy liability calculations should be anchored in the expected costs based on the actual position being held/expected to be held, which implies using a P-measure approach applied to the net exposure. Therefore, this is the basis that the task force is recommending in this paper.

Where hedging strategies are being employed to help mitigate risk, the net exposure itself should reflect the risk mitigation and costs of the hedging strategies. Determination of the costs of hedges should normally be determined using a capital markets framework, even though the P-measure basis applies to measuring the overall risk exposure.

The model should not generate negative stock prices or negative interest rates.

State dependent models relate the change from one period to the next to current market levels or recent market performance. For example, a mean-reverting process is state dependent because the future scenarios depend on how the current market variables relate to long-term historical values. State dependent models are not required, but are acceptable if they are justifiable based on the historical data and meet the calibration criteria.

A related issue that receives a significant amount of discussion is whether the model should explicitly allow dampening of the impacts of recent market experience (e.g., reflect an assumption that following significant appreciation, a higher provision for a correction is appropriate and vice versa). This is another form of a state dependent model so such behavior assumptions are permitted provided they continue to meet the calibration criteria.

e) *Stochastic Model Parameters Estimation*

Different models may require more or less parameters and refer to different statistical distributions. A typical model should at least have two parameters relating to the drift and volatility of the stochastic process.

These model parameters should be estimated based on historical market data as opposed to recent market performance. Due to the long-term nature of the segregated fund guarantee, as a rule of thumb, historical data should cover at least two times the projection span. However, when historical data are not available or it is not justifiable to use it, then some adjustments may be required.

Generally, market indices should be modelled rather than the specific fund performance. There will be more credible data available for the market index and the specific fund performance can depend on additional factors that may not be consistent over time (for example the fund manager can quit or be replaced).

Parameter estimates for a number of different market indices may need to be included in the generated scenarios so they can be combined to model a specific segregated fund portfolio. When more than one index needs to be projected, it is necessary to allow for correlations between different markets. It is not necessary to assume that all markets are perfectly positively correlated, but it would be appropriate to use correlations other than zero. The actuary should consider that correlations are not stationary, and that they tend to increase during times of high volatility or negative returns.

If making *ad hoc* adjustments to observed correlations, care should be taken to ensure that the resulting correlation matrix is internally consistent. (Technically, a correlation matrix should be positive semi-definite).

Also, when foreign indices are used to establish the benchmark index, the foreign exchange rate must also be considered. In some situations, it may be appropriate to have separate parameters for the market index and for the foreign exchange rate. The fact that a currency has depreciated or appreciated significantly in the historical period should be scrutinized before assuming that the trend will continue in the future.

If required, these parameters must be adjusted to reflect the skewness and the tail fatness observed in the historical data. This required adjustment is discussed below as part of the calibration process.

The model parameters are not required to be constant over the projection horizon.

f) *Selecting Investment Return Assumptions for Specific Funds*

To develop investment return paths for a specific fund, an appropriate proxy for the segregated fund must be constructed. The specific fund's investment policy, its asset allocation implied by the fund performance objective, the history of fund performance and trading activities must be examined prior to proxy construction and then reflected in the proxy asset composition.

The proxy may take the form of a linear combination of recognized market indices or economic sector sub-indices or, less commonly, as a well-defined set of trading rules in a specified asset universe. Using combinations of recognized market indices or economic sector sub-indices facilitates using a limited number of well developed and researched data-sets to model a wide range of funds.

The proxy construction process should involve analyses that confirm a close relationship between the investment return proxy and the specific segregated funds.

The specific analyses can include, but are not limited to:

- Comparison of the serial long-term and short-term historical returns of the proxy and the specific fund.
- Analysis of serial correlations between the proxy and the specific fund.
- Comparison of asset composition over time of the proxy and the specific fund.
- Comparison of the systematic risk between the proxy and the specific funds' assets.
- Comparison of the specific risks between the proxy and the specific funds' assets.
- Comparison of the source-of-return attribution between the proxy and the specific fund.
- Comparison of the volatility and risk-adjusted return between the proxy and the specific fund.
- Comparison of the long-term expected asset composition of the proxy and the specific funds.

When sufficient historical information about specific funds' performance is not available, the proxy should be constructed by combining asset classes and/or allocation rules that most closely reflect the expected long-term asset composition of the specific fund. The proxy return-generating process can then be modelled by mapping this asset composition to the historical performance of market indices or economic sectors that most closely reflect the proxy long-term asset composition. Where sufficient historical information for a specific market index or sub-sector does not exist, the return-generating process would reflect the contribution of this component to the specific fund total return by reference to the efficient markets risk-return relationship, as described below.

Investment managers may seek to generate incremental returns by short-term changes in fund allocation to individual assets or asset classes/sectors. As described below, such incremental returns may only be achieved at an increased level of risk. This risk component must be reflected in the return-generating process of the specific fund.

A well-established tenet of the modern portfolio theory is that, over the long term, additional return can only be achieved by undertaking additional risk. If the specific fund investment policy expects to generate excess return by pursuing active portfolio management, a risk-return relationship must be reflected in the specific fund's return-generating process. This relationship can be captured from efficient frontier construction, the capital market pricing model or arbitrage pricing theory. The final proxy for the return-generating process of the specific fund should conform to this risk-return relationship.

2.1.2 Calibration of Investment Return Models Used for Generating Returns

The calibration tests are to ensure that the model is able to generate scenarios that take into account the tail skewness and/or fatness observed in historical data. The emphasis of these tests is placed on fitting the tail of the distribution as opposed to fitting the entire data set or some other measure such as the mean.

Calibration requirements are included only for equity return models, since this is the primary source of risk with respect to segregated fund investment guarantees.

For equity return models, the model should be calibrated using a prescribed data set. The recommended data set is the TSE 300 Total Return monthly data from 1956 to 1999. Once the model has been calibrated with this data set, the "fitted" model should be used for all indices as described below (in other words, the basic model is only "fitted" once).

(continued on page 12)

Modelling
continued from page 11

For models which are a function of recent history (e.g., market levels, current volatility, mean-reversion process, etc.), calibration tests must be done using the long-term trend of these parameters as recent history. In other words, the model calibration should be done using long-term trends in values for the recent history, and not use the actual history of only the past immediate periods. Once the model is calibrated, the forward projection from today can start with the actual values for the immediate past periods. The task force proposes the following prescribed calibration process for stochastic models of total returns on equity investments.

1. All model calibration for equity return models should be done with a single data-set. The data-set we are proposing is the TSE 300 total return data from January 1956 to December 1999 (end-of-month values). The parameters should be estimated by maximum likelihood techniques or by similar statistical methods. No allowance should be made for expenses in the parameters.

2. The calibration is applied to the total return one-year, five-year and ten-year accumulation factors generated by the asset model. For models which assume correlation between returns in successive periods, the accumulation factors should be calculated using neutral starting values.

3. Table 1 provides maximum returns for the 2.5th, 5th, and 10th percentiles for the accumulation factors (Appendix C provides a description of the analysis undertaken to establish these calibration points). As an example of how to interpret the table, for a five-year holding period, the total return must be -25% or lower at least 2.5% of the time.

Accumulation period	2.5 percentile	5.0 percentile	10 percentile
one-year	0.76	0.82	0.90
five-year	0.75	0.85	1.05
ten-year	0.85	1.05	1.35

4. The model with the initially determined parameters (i.e., uncalibrated parameters) might not satisfy the calibration criteria in Table 1. In this case the parameters may be adjusted until a set of calibrated parameters that meet the calibration criteria are determined. Alternatively, a different model may be selected.

5. The final calibrated parameters for the TSE data-set should be extrapolated to other data-sets using the formula that follows. If k(TSE) is the uncalibrated parameter for the TSE data-set, and k1(TSE) is the calibrated parameter, then for any other data-set, the calibrated parameter k1(DATASET) is defined as k1(DATASET) = k(DATASET) + [k1(TSE) – k(TSE)]. This approach should be followed for each fitted parameter.

6. Each of the maximum return criteria must be met. This means that the model used must produce return values for the accumulation factors that are no larger than the appropriate table values, for each holding period/percentile combination.

7. For some models the percentiles may be calculated analytically; if simulation is used care must be taken to avoid bias in the random number generator. A sufficient number of simulations should be performed to ensure that the criteria are met with a high degree of confidence (95% certainty would not be unreasonable).

8. In addition to the percentile criteria in Table 1, the mean of the one-year accumulation factor should lie in the range 1.10 to 1.12. The standard deviation of the annual accumulation factor should be at least 0.175.

Appendix A provides an example of how a common simple fixed volatility lognormal model can be calibrated to meet these criteria.

Other models are equally acceptable, and indeed may be preferable if they do a better job of capturing the characteristics of actual market returns (such as fat tails and time varying volatility). Appendix B provides a brief overview and further references for how other models may be calibrated (e.g., regime switching lognormal, stochastic volatility lognormal, stable model).