Office of the Policyholder Advocate

Appendix 14

Approaches to Modelling:

Stress Testing, Scenario Analysis and Stochastic Modelling

A report for the policyholder advocate in connection
with the reattribution of the inherited estates
of the CGNU Life and CULAC with-profits funds

June 2009

Report By: P. Booth BA, FIA, FSS, PhD**,** Professor of Insurance and Risk Management, Faculty of Actuarial Science and Insurance, Cass Business School

Contents

1.00 Introduction

1.01. Preamble

In this Appendix we will examine different approaches to understanding and managing the risk in a financial institution – with particular reference to a life insurance company. The approaches we shall examine fall into two categories: 'deterministic' and 'stochastic'.

A deterministic approach typically involves choosing particular values for financial variables and projecting the financial position of the company going forwards. Deterministic approaches can be used to project the financial position of the company under a range of different assumptions. However, when deterministic approaches are used there is no direct attempt to assess the probability of different outcomes occurring[1].

Stochastic approaches, on the other hand, use statistical models to try to determine the probabilities of different financial outcomes. They reflect directly the fact that the various financial variables that affect the position of a life company are 'random variables' (i.e. we cannot predict their outturn precisely but we can try to estimate what values they will take with particular probabilities). Stochastic statistical models use, as inputs, the (sometimes quite sketchy) information that we have about the probability that different financial and non-financial variables will take different values and the statistical relationships between different financial variables. Using such inputs, a stochastic model can be used to make a series of statistical projections which, in theory, can estimate, under different starting assumptions (such as assumptions for the allocation of assets between different investment classes), the probability distribution of financial outcomes for a financial institution.

Stress testing and scenario analysis are two forms of deterministic approach to understanding risk in a financial institution. The Financial Services Authority (FSA) often talks about stress testing as being a process by which the financial position of a company is recalculated after allowing for a change in a given parameter (e.g. a fall in equity values). Scenario analysis is often intended to imply a process of calculating the change in the financial position of a

[1] The word, 'direct' is important here: see later discussion.

company after a change in a range of parameters. Scenario analysis can be 'dynamic' (as in a process known as 'dynamic solvency testing' which is frequently used in life insurance companies) in that the parameters can be allowed to change over time and, perhaps, one parameter can be allowed to change in response to another changing (e.g. bonus rates are reduced when investment returns fall). It is generally accepted that, in scenario analysis, there should be internal consistency between the values chosen for connected variables. For example, if it is desired to project the financial position of a company in times of very high inflation, it might be expected that expenses and nominal long-term interest rates will also rise and be assumed to be higher. Having said that, each scenario must be justified on its own merits. Thus, for example, if it is regarded as plausible that nominal long-term interest rates could be low at a time of high inflation, then it may be desired to test the financial position of the company against this eventuality.

For much of this Appendix, we shall not distinguish between different deterministic approaches of assessing the risks facing a life office as much of the discussion is about the merits of deterministic and stochastic approaches in general. The key point is that commonly used approaches such as stress testing, scenario testing and dynamic solvency testing are all processes by which the financial strength of an insurer is tested against *deterministic* changes in one or more parameters and that these changes can take place once and for all or can follow a process through time (for example inflation may be assumed to increase by 1% per annum in each of the next ten years).

Stochastic modelling, on the other hand, refers to the process of determining a probability distribution of the financial position of the insurer at a given point(s) in time, given a statistical model for the various relevant parameters and for their interactions.

Where stress testing is used, the chosen scenarios should have some rigorous justification. They could, in fact, themselves be the outcome of a statistical modelling process. A statistical model can be used to determine (say) the joint probability of a rise in gilt yields by 2% and a fall in equity markets by 20%. The directors can then decide, given that probability, whether they wish to use this as a scenario for the purpose of stress testing and so on. The scenario would then be used in a deterministic way to determine the change in the solvency position of the insurer.

The main part of this paper examines the role of stochastic and other approaches in setting capital for and assessing risks facing insurance companies; the appendix to this paper examines the role of stochastic models and other approaches in the valuation of liabilities.

1.02. In Anticipation of the Conclusion…

In theory, stochastic models give the user the precise output that is needed to assess the risks faced by a financial company. The user can calculate, under different assumptions for parameters such as the make-up of the investment portfolio, the probability distribution of future outcomes and hence the probability of failure. Stress testing cannot do this. Stress testing can show the outcome when particular financial events occur. The problem with stochastic modelling is the difficulty we have determining whether the probability distribution of outcomes is at all realistic. With stress testing, one can check whether the individual scenarios used are appropriate, but one is limited in the range of variables that can be processed and modelled. Both procedures should be useful in alerting management to the risks that a company faces (particularly non-linear risks[2] such as those involving options) but I am sceptical that either can accurately determine how much capital needs to be held to (say) limit the probability of insolvency to 0.5%. Some years ago I came to the conclusion that stochastic models should be used for understanding business risks but we should be careful how reliant we are on their results. With regard to decisions such as those surrounding how much capital should be held, they should not be assumed to come up with a single "correct" answer to the question that is asked of them. The approach used for building stochastic models is necessarily pragmatic and itself judgmental.

[2] Some risks are what might be termed 'linear': for example, in many situations, if equity returns fall, the financial position of the company will worsen with all lower values of equity returns producing a worse outcome than higher values. With non-linear risks the outcome does not change proportionately with changes in underlying variables. Bankers are often said to face non-linear risks with regard to their bonus packages. If their performance improves their bonus gets bigger. Up to a point, if their performance worsens, their bonus gets smaller. However, if their performance is catastrophic, their bonus cannot be reduced below zero (so there comes a point where a banker's bonus stops getting smaller as performance worsens). There are a number of non-linear risks in life insurance companies.

2.00 Some History and Context

2.01. Stress Testing

Life offices have calculated their financial positions using a range of different assumptions for many decades (for example, we used about ten different bases at Equity and Law in 1986 when I worked there). Formal stress tests with some statutory backing started just over 20 years ago. Life offices were required to demonstrate that their balance sheets could withstand particular movements in asset markets (and by implication rates of interest used for valuing liabilities). This led to companies being required to hold a so-called resilience reserve. The resilience reserve scenarios were, rightly, criticised for combining together movements in equity values and gilt yields that were highly unlikely ever to happen simultaneously. Thus the joint probability of some of the scenarios happening was negligible. A more realistic (though only just intelligible) variant of those scenarios is used today in the FSA's twin peaks approach to statutory solvency.

In Canada, the statutory solvency tests require companies to conduct 'dynamic solvency testing' which is essentially a dynamic version of stress testing that allows assumptions relating to financial and non-financial variables to change through time (see below). It is a purely deterministic process. In the UK, there is nothing to stop a life insurance company using any of these techniques, or variants of them, to determine how much capital to hold – as long as the resultant capital is not less than that which the FSA requires the company to hold. This approach of using simple dynamic solvency testing to determine how much capital to hold would generally be sufficient to stop serious financial risks leading to insolvency. For example, it can easily be used to examine what happens to liabilities with non-linear payoffs[3] as financial and non-financial assumptions change. There is a recent discussion of the use of dynamic solvency testing in Shiu et al (2006).

[3] A good example of a policy type with a non-linear payoff is a guaranteed annuity option. The precise nature of the payoff will vary from policy type to policy type, but one possibility is a situation where the value of the liability is constant if long-term interest rates should rise from current levels but rises as long-term interest rates fall, once they have fallen below a particular threshold.

2.02. Stochastic Modelling

Stochastic modelling has developed over the last 20 years or so. It was originally developed to price maturity guarantees and can, in many circumstances, provide a good "feel" for the probability of a company having to pay out more than a certain percentage of asset shares. However, the more complex the applications, the more apparently precise the outputs and the more complex the models, the greater are the difficulties of really understanding what is going on.

It is worth contrasting the approach used for the calculation of value-at-risk in banking groups with that used for stochastic modelling in life insurance. Both have developed pretty much simultaneously though largely independently. In value-at-risk modelling a very short (often five day) time horizon is used to project the value of a bank's assets and liabilities forward. Daily data over a few years can be used to derive the models and the parameters (often with no economic, just statistical, input). Even then, instability in the data is a problem that is recognised by value-at-risk model users, as are the dangers of non-linearities and fat tails[4] that may not be represented by the fitted probability distributions. Indeed, with value-at-risk modelling, at least there is some data to try to ascertain how fat the tails should be so that adjustments to the results produced by standard models can be made. With stochastic modelling in a life company, one is trying to project several decades forward and yet the models are built using a data period that is much shorter than the projection period. This is probably a unique situation in financial modelling (though Nicholas Stern adopted a similar approach in his recent modelling of the economics of climate change). Stochastic investment modelling in a life insurance company may be useful as a management tool and it may help us understand the risks the company faces and the capital it would be useful to have given those risks, but we should be wary of those who try to make the models work too hard.

[4] It is often helpful to represent the probability distribution for variables as being symmetrical (or being some function of a symmetrical distribution). However, it is often thought that investment returns from portfolios may be 'fat tailed' because, amongst other reasons, in difficult times, markets might be more likely to 'collapse together' than is implicitly assumed by the model structures that are often used. 'Fat tailed' refers to the fact that there is a higher probability of very adverse outcomes than is assumed in simple model structures. In theory, models can be adapted to deal with this problem but truly catastrophic events might be genuinely different from previous events and thus not be reflected in the dataset from which models are designed or parameterised.

2.03. Stress Testing for Regulatory and Internal Management Purposes

As has been noted above, stress testing is required for certain statutory purposes. For a number of years, insurance companies were required to hold a “resilience reserve” which was intended as a capital buffer and was calculated so that the insurance company had sufficient capital to withstand given changes to financial variables (specifically changes in gilt yields and equity values, in the early years of the test). The resilience tests were widely criticised over the years and it was felt that they could lead to perverse results; as such they were frequently changed. The tests have since evolved and become more complex and they are used in the ‘twin peaks’ approach to capital assessment for the major with-profits life insurers subject to the ‘realistic reporting’ regime that the FSA introduced at the end of 2004. Specific stress events, proposed by the FSA, are used to determine the capital that needs to be held under both the realistic and the statutory reserving basis peaks. The stress tests used are similar in both cases, though, in the realistic peak calculation, scenarios relating to persistency and credit spreads are also included.

Companies are also required by FSA rules to undertake an “Individual Capital Assessment” (ICA) using their own models. This can use stochastic modelling, or stress testing, or both, although stochastic modelling is the norm, and indeed with-profits life insurers would struggle to do the assessment without it. The purpose of the ICA is to determine the amount of capital that it is necessary to hold so that the company has a probability of insolvency over the following 12 months of 0.5% or less. Although the models that are used in this calculation are not specified by the FSA, if a company produces an ICA that the FSA believes to be too low, the FSA can impose its own capital requirement through the issue of ‘Individual Capital Guidance’. In a sense, this would be an implicit disapproval, by the FSA, of a company’s internal models or the parameters in the model.

For internal management purposes, companies could then use both stress testing and stochastic modelling to determine the amount of capital that they wish to hold, taking into account both economic and regulatory considerations. There would, of course, be no restrictions on the forms of the models used for this purpose.

3.00 Stochastic Modelling: Further Analysis

Stochastic modelling has been used for various different purposes. As has been noted it was originally developed to price and assess the risks of maturity guarantees. It was then quickly developed for other purposes.

In both pensions and life assurance, stochastic modelling can be used to assess the risk and return profile of different investment strategies. The correlations between the assets and liabilities is particularly important in assessing the appropriateness of different investment strategies, as is the possibility of management reaction to particularly adverse scenarios (for example by reducing bonuses or changing the emphasis between terminal and reversionary bonus). Indeed, an important part of the modelling process is that it will help actuaries and directors obtain a better understanding of the relationships between different asset portfolios and the liabilities and the actions that can be taken to reduce risk.

In life assurance, stochastic modelling can also be used to determine the pattern of bonus that is achievable with different investment strategies – amongst other things modelling can help the company determine the investment policy that will produce a bonus pattern that is compatible with 'policyholders' reasonable expectations' as well as the impact of different smoothing strategies. Models can also be used to assess the risks from new business and the costs of new business as well as to determine the pricing structure of new business.

The applications in the two paragraphs above use the general shape of the distribution of outcomes in the decision-making process. Both the mean outcomes and various measures of dispersion can be used in the analysis. One hypothetical example, of how stochastic modelling could be used would be to determine the implications for the variance of outcomes from diversifying the equity portfolio of an insurer into a range of overseas markets.

Stochastic modelling is also used to determine the ruin probability (probability of insolvency) both on an economic basis and on a statutory basis for a given set of input parameters (including the investment strategy and level of capital). Alternatively, it can be used to determine the level of capital and/or investment strategy to achieve a particular probability of insolvency. Thus, we can think of the model being used to determine the right level of a number of control variables (such as the level of capital, ruin probability or investment mix) in order to achieve a particular value for the other variable (e.g. for a given investment strategy and ruin probability we can use the model to determine the level of capital that has to be held).

The problem with this type of application is that it is the tails of the probability distribution that are being used. The probability of statutory insolvency derived from stochastic models tends to be very small for most companies. The probability of economic insolvency (i.e. a firm having a negative net present value) is even smaller – and economic insolvency is as likely to happen as a result of operational risks that cannot be modelled than as a result of the kind of risks that are captured by stochastic models. There are several problems with using the tails of stochastic models – we are constructing models for use over long time periods when extreme events have happened very infrequently in history. More generally, it is well recognised that the tails of the distribution of outcomes in life insurance stochastic models are poorly understood. Whether a particular level of capital gives rise to a 1% or a 2% probability of insolvency might be extremely difficult to determine in practice, never mind whether it gives rise to a 0.1% or 0.2% probability of insolvency. Thus a crucial message of this paper is that stochastic modelling should be used to understand the risks of the business, and we should be wary of making definitive statements about the probability of particular outcomes under different assumptions. It is worth noting a comment by Paul Sweeting (Chairman of the Faculty and Institute of Actuaries' Finance and Investment Practice Executive Committee) in the editorial of the actuarial profession's magazine, "The Actuary" (January/February 2009 edition). He comments: "While modelling can be used to give an indication of the level of risk being faced, it is vital to recognise that it is only an indication, best used when combined with subjective checks. Models can be unreliable just when they are needed most, particularly in the tails of the distribution."

As both stress testing and stochastic modelling can be used to help an insurance company determine how much capital to hold, both techniques can be used both for internal management purposes and for regulatory purposes. As has been noted, and as will be discussed further, insurance companies have always calculated their liabilities on a number of bases – implicitly at least using stress tests to determine their capital levels. Statutory valuation requirements then evolved to require an insurance company to hold a capital layer determined by a stress test (resilience test as it was known) – the stress test for this particular purpose was determined by the Government Actuary's Department (now by the FSA). The FSA still applies such a test although the practical details have changed somewhat. Stress tests are still also used for many management purposes (including deciding how much capital to hold) with the scenarios determined by the company. Stochastic models can be used for some aspects of the statutory valuation process and here the FSA requires models to have certain features. However, in general, stochastic models are designed and calibrated by users themselves (and by their consultants). Indeed, it is a developing field with models continually being developed and refined (not least for reasons discussed elsewhere in this note).

4.00 Stress Testing versus Stochastic Modelling

4.01. General Discussion

Although intrinsically deterministic, we can still relate the results of stress testing to probability statements. In particular, the impact of particular events can be quantified and then the likelihood of those events examined.

Stress testing is often contrasted with value-at-risk modelling in the securities markets literature. As has been noted above, value-at-risk modelling can be regarded as the banking equivalent of stochastic investment modelling, so these comparisons in the literature are relevant to this discussion. Stress testing focuses on estimating the financial losses that would arise from particular extreme events. When interpreting the results of a stress test we can try to impute a probability of occurrence to a stress-test event from available data and prior theory. Stochastic modelling, on the other hand, requires the direct use of statistical modelling. For a given asset portfolio (and set of liabilities) a stochastic model determines a probability of losing a particular amount of capital, without necessarily explicitly defining the events that will give rise to the loss. The calculated probability of losing a given sum will be a function of the particular statistical model used and the results need to be interpreted in the light of the characteristics of the model.

Shaw (1997) suggests that there are significant shortcomings of stress testing. Specifically, he suggests that there are two obvious questions: where do the scenarios come from? And, what should one do with the results? He argues that stress testing does not add value if realistic estimates of the probabilities of the combinations of events happening cannot be made: if we have the models to make those estimates, then why not use those models for value-at-risk modelling in the first place he then asks? Shaw suggests that value-at-risk approaches generate the “correct numbers” for analysing risk because they take into account the joint probability distribution of the occurrence of possible events. Further, Shaw proposes that the value-at-risk approach should be developed to overcome its shortcomings, rather than combining current value-at-risk techniques with stress testing. Shaw is really saying that stress testing can add nothing that cannot be provided by statistical modelling.

At a theoretical level, the criticisms by Shaw are reasonable. Assuming we know the joint probability of all possible adverse events happening, there is no information that a stress test provides that value-at-risk modelling does not provide. However, in normal business life, we do not have all the information about the distributions of investment returns that we would like – and this is particularly the case for the long period over which we model life companies’ financial positions.

Stress testing can be put into a statistical framework if necessary. A stress test will model the amount that will be lost from particular events. As has been noted, it may then be possible to ascribe a probability to these events or at least possible to acquire some useful information about their likelihood using statistical models.

Though, Shaw’s assertion is thus correct at one level, in any risk modelling the estimated probabilities in the tails of the probability distribution may be incorrect because of the infrequency of the outcomes which make up the tails and because of the breakdown of the “usual” correlations between events in times of financial crisis. In addition, the probability distributions cannot always be expected to be stable over time: regime changes, learning effects and so on, will lead to changes in the probability of and effects of extreme events. In any of these circumstances the results from a statistical risk model may be difficult to interpret or the model may be an insufficiently accurate representation of reality. The more complex the portfolio and the more scarce the data, the more difficult it will be to develop a reliable statistical model.

Starting, as stress testing does, from the perspective of asking what events lead to problems for the company and then uncovering information about the likelihood of those events can make the risk management and capital setting process less opaque. In general, in such an approach, the subjective assessment of probabilities of extreme events is necessary. This can be easier to handle in a stress-testing framework because we are first examining the events that will lead to financial problems and then focusing on determining the probability of those specific events. This does not rely on a statistical model being "correct" in general as we are only using it to ascertain the probability of some particular events. On the other hand, we cannot use this approach to determine the probability of all possible events that threaten the company.

If stress testing and statistical risk modelling are both carried out, it is important to check that their results are consistent. For example, if a stochastic model suggests that £1bn of capital is necessary to reduce the probability of failure to 0.5% or less, yet a stress test concludes that an event to which the managers of the institution subjectively attach a higher probability than 0.5% would lead to (say) the loss of £2bn of capital, the model and the assessment of the likelihood of the stress event should be investigated further.

A further advantage of stress testing is that non-market risks can be factored into the computation of the impact of the stress events (e.g. changes in regulation that lead to increasing costs).

It should be mentioned, at this stage, that a further issue is the relative merits of stochastic modelling and other techniques for assessing the *expected* value of a balance sheet variable or policy benefit (such as the liabilities of an insurance company, the value of the inherited estate, or expected payouts to policyholders pre and post reattribution in the case of this transaction). We have concentrated in the sections above on the use of techniques to understand risk in relation to events which might be expected to occur only infrequently. In the Annexe, we examine the relative merits of stochastic and deterministic techniques to determine the *expected* values of various financial outcomes.

4.02. Advantages of Stochastic Modelling: a summary

- The points made by Shaw are theoretically correct. More information can be incorporated into a stochastic model than into a stress testing model. Theoretically, all relationships between relevant variables can be incorporated in a stochastic model.
- The output of a stochastic model (e.g. a probability of ruin for a given level of capital) is precisely the output that we require to make judgments about how much capital we need to hold etc.

4.03. Advantages of Stress Testing: a summary

- We cannot assume that there is no limit on the discovery and application of knowledge regarding the joint probability of events. In other words, the estimates that come from stochastic modelling must have a standard error attached – but do we even have the information to calculate the standard errors given the uncertainty about whether the stochastic model is the right model? There are particular problems with the reliability of the tails and the behaviour of correlations when extreme events happen. Recent events in financial markets would seem to have provided more evidence for this.
- Stress testing can focus on the particular events which give rise to losses and the probability of those events can be analysed explicitly. A stochastic model tends to lose information about specific events in the context of more general relationships. In particular, stress testing can be useful in focusing on the effect of particular non-market risks and operational risks (e.g. costs of one-off systems failure and/or changes in mortality).

- There are many subjective elements in stochastic modelling and the subjective and controversial aspects of a model are not easy for an outsider to define and explicitly discuss. It is also not easy for an outsider to determine the impact of using one set of parameters or one modelling technique rather than another.
- The simplicity of stress testing is also probably better suited to "learning from experience". If particularly damaging events happen that were not chosen as one of the stress scenarios, it is easy for the Board to ask why they were not chosen and for discussion to ensue about improving the process. It is much harder for a Board to have feedback on the results of stochastic modelling – or, indeed, even understand why a stochastic model might appear to understate or overstate the probability of extreme events in retrospect.

4.04. Conclusion

The approaches of stochastic modelling and stress testing are complementary. In a well-managed financial institution, one would expect to see both used. Whilst directors would wish to see numbers such as the amount of capital that has to be held to limit the probability of insolvency to a particular level, the complex nature of stochastic models is such that they would find it very hard to interpret the information or question the results: just as a policyholder advocate will find it difficult to question the results of stochastic modelling. I would expect that directors would also wish to have information about the impact of specific events on the capital of a company through a process of stress testing. My concern is not with the use of either or both of these methods but I question the idea that either method can be used to make a precise statement about the amount of capital that needs to be held to limit the probability of ruin to a particular level. In the words of Professor Steve Haberman in a sessional meeting of the Institute of Actuaries: "The [stochastic] model provides a self-contained and consistent way of generating simulations or scenarios. The estimates of moments, percentiles, etc. cannot be regarded as unbiased, but they can be informative as general indications of results, especially if conducted alongside a thorough sensitivity analysis and monitoring of the results as they emerge." That statement was made ten years ago but I still believe it holds true[5]. Indeed, recent events in financial markets have surely confirmed this way of thinking, which sits comfortably alongside the statement made by Paul Sweeting (above). Even where models were used in an investment banking context, with more data and shorter time horizons, it seemed that they were unable to capture properly the nature of the risks in some securities markets – especially, though not only, when liquidity dried up. The way in which regulators encouraged institutions to have similar approaches to modelling may well have contributed to the failure of the models as the impact on financial markets of the similar behaviour of different institutions might not have been captured by the modelling process. All these points should make us wary of using the tails of a stochastic model as if they were providing precise measures of the probabilities of particular outcomes.

5.00 Stress Testing: a framework for analysis

5.01. The Process of Stress Testing

There is a helpful statement in a recent Bank of England Financial Stability Review (Bank of England, 2007, page 11) of how stress scenarios should be chosen. It suggests the use of "severe, but plausible, adverse future stress scenarios" and comments, "these scenarios need to take adequate account of potential amplification and feedback effects that might occur in highly connected markets in times of stress, including those arising as a result of a fall in market liquidity." The Bank of England suggests that problems such as those caused by the collective behaviour of a number of firms should be taken into account. An example of this sort of behaviour in the life insurance field occurred in late 1998 and could occur again. In 1998 solvency problems caused life insurance companies to buy long-dated gilts to improve the match with their liabilities. This lowered long-dated gilt yields, thus causing solvency problems for other companies that were more heavily invested in equities (and further problems for companies that, despite their efforts to buy more long-dated gilts, still had liabilities with a longer duration than their assets).[6] Companies then tried to match their liabilities better by further selling of equities and buying of bonds, thus lowering equity

[5] It is also a fair summary of F. A. Hayek's view of econometric modelling (see, for example, Hayek, 1988).

[6] The problems arose because, when interest rates fell, the present value of life insurance companies' liabilities increased. This was not offset by increases in asset values, even where insurance companies had invested in bonds, because the durations of fixed-interest bond portfolios were generally less than that of life insurance companies' liabilities.

values and bond yields further and worsening the problem. It is worth mentioning that the long-dated bond market is now much more liquid than in 1998.

Page 44 of Issue 21 of the Bank of England's Financial Stability Review suggests a useful systematic way of applying stress testing[7]. The approach has been developed for analysing risks across the banking sector; however, there is no reason why it could not be applied to a specific life insurance company. The four proposed steps (with the potential life insurance counterparts or examples in brackets) are:

(1) Identification of the potential vulnerabilities of a bank (identification of the vulnerabilities of a life company e.g. to falls in equity markets, falls in fixed-interest yields etc.).

(2) The construction of stress scenarios that expose the vulnerabilities (specific falls in equity markets and/or gilt yields).

(3) A mapping of the channels through which the scenarios will affect the economy (channels through which the scenarios affect the life company – e.g. falls in asset values leads to a rise in costs of guarantees, embedded options and so on. These channels should take account of management and policyholder actions and actions by other companies that could either mitigate or exacerbate the impact of the stress event).

(4) An estimation of the impact on major banks (the estimation of the impact on the balance sheet of a life insurance company).

5.02. Stress Testing: Choosing Scenarios

To these four processes, it can be added that there should be an analysis of the likelihood of combinations of the chosen events used in a stress test. Furthermore, the likelihood of the events should not just be considered individually, the combination of events should be seen as a consistent scenario. For example, a large fall in gilt yields combined with a fall in equity

[7] There is more detail provided in the July 2006 Bank of England Financial Stability Review (Bank of England, 2006) and in Haldane et al (2006).

values is relatively unusual unless there is a large fall in inflation expectations (as in 1998). Indeed, it was the use of a combination of events that were individually reasonably likely to occur but together where very unlikely to occur that brought the original UK life insurance company statutory resilience test into disrepute. There is no reason why stochastic models (and other methods of data analysis) cannot be used to determine the likelihood of stress events and combinations of events.[8] It should be noted that, even if the probability of occurrence of a particular event does not change much over time (say the probability of a fall in the equity market by 25% in one year), the probability of occurrence of potentially linked events (such as a simultaneous change up or down in gilt yields) may change as a result of changes to fundamental structures of the economy or changes in the way investors are likely to react to a particular event. Stochastic modelling may not capture such changes effectively because stochastic models do not focus on the data relating to *specific* combinations of events.

Thus, in conclusion, as far as insurance company modelling is concerned, the scenarios chosen should have an important impact on the insurance company (i.e. they should be *practically relevant*) and, where a group of events rather than a single event is being considered, the different aspects of the scenarios should be mutually consistent and their likelihood considered in a coherent financial and economic framework. The scenario should include the reaction of management and policyholders and, where appropriate, possible reactions of other companies which may either exacerbate or ameliorate the severity of the stress event.

6.00 Dynamic solvency testing

This part of the note describes the process of dynamic solvency testing (DST) in life insurance. This is simply a dynamic form of stress testing designed to ensure that the business has sufficient capital to protect it from particular economic, financial and non-financial events.

---

[8] However, it does not, of course, follow that when an event of (say) probability of occurrence 0.05 gives rise to a loss of £500m, that the 5% value-at-risk is £500m.

Dynamic solvency testing can be used as a tool by regulators to calculate regulatory capital, a tool by insurance companies to help them understand the risk to the company, or as a tool by insurance companies to help them determine economic capital. A recent paper by Shiu et al (2006) discusses the practical application of dynamic solvency testing in the UK life insurance industry.

6.01. Introduction to dynamic solvency testing

Using dynamic solvency testing, we can:

- Model the impact of changes in financial conditions forward over several time periods.
- Model the impact of changes in financial conditions which evolve over several time periods.
- Model changes in business conditions (for example, changes in new business sales or in the profitability of new business).
- Model serially dependent changes in financial conditions (for example, if it is believed that markets follow cycles).
- Model the impact of changes in financial conditions on the profit and loss account and on cash flows.
- Model the impact of management reaction to changes in financial conditions (for example, if interest rates fall and the solvency position is impaired, an insurer may change its investment stance to try to immunise itself against further interest rate changes).
- Model the risks of policies with options: with regard to guaranteed annuity options, for example, it is possible for dynamic solvency testing to test the effects of changing interest rates and mortality rates and different patterns of consumer behaviour in response to such changes.

Dynamic solvency testing can be distinguished from stochastic investment modelling in that the scenarios chosen are deterministic: there is no direct attempt to find the probability distribution of potential outcomes for the business. In this respect it is simply a dynamic version of stress testing. It should be noted that, in the insurance literature, the phrase "dynamic financial analysis" (DFA) is used to describe both deterministic and stochastic methods of financial projection.

6.02. Objectives of Dynamic Solvency Testing

The objectives of dynamic solvency testing are described by the Canadian Institute of Actuaries as, "The process of analysing and projecting the trends of a company's capital (solvency) position given its current circumstances, its recent past, and its intended business plan under a variety of future scenarios". The principal goal of the process is, "To help prevent insolvency by arming the company with the best information on the course of events that may lead to capital depletion, and the relative effectiveness of alternative corrective actions." It is stressed that it is the degree as well as the timing of any capital depletion that is important.

In both the UK and Canada, dynamic solvency testing is generally tied to statutory valuations. It is used to provide the insurer with an indication of the events that would cause it to become insolvent on a statutory basis. Dynamic solvency testing can be used to assess the circumstances in which an insurer will fall below any pre-determined solvency level and thus assess the margin above the required regulatory solvency margin that the insurer will want to hold. However, dynamic solvency testing does not have to be tied to statutory valuations in this way. It can also be used to assess the financial impact on the underlying economic solvency position (i.e. realistic solvency position) of an insurance company. Thus it can be used to help an insurance company set its levels of economic capital.

7.00 The Construction of Stochastic models

7.01. Features of Good Stochastic Models

If we are to use a process of stochastic modelling, then we have to consider the appropriate features of any model. All models are a simplification of reality, but what features should a good asset and liability model have?

- It will be an effective representation of reality. It should mimic the real-world features of asset performance and the interaction of different asset classes with each other and with the liabilities. This consistency of the treatment of assets and liabilities is essential.
- It should be economically intuitive. Not all economists agree on the way in which different variables interact but, for example, if we are modelling short-term interest rates, we should expect some relationship between short-term interest rates and inflation. Similarly, in a pension fund liability model, we would expect a relationship between nominal wages and inflation.
- Models should reflect the historical data.

7.02. Stochastic Models and the Usual Processes of Scientific Investigation

There is a fundamental problem with stochastic investment models designed for long-term simulation: the models developed cannot be subject to the normal scientific process of falsification due to the time period over which the data is gathered relative to the time period over which the modelling takes place (see also Section 3). For some markets, a single 20-year data set is used, for example, to build a model that is designed to project simulations forward over four or five decades. Certain processes can be tested and incorporated in the model (for example the assumption of efficient markets) and because the processes are assumed to be invariant with time, this is not a particular problem. However, this is not the case with some of the most important features of stochastic models. For instance, a particularly important feature of many models is the relationship between various variables and inflation. This includes both the relationship between asset classes and inflation (such as real estate rents and real estate values) and the relationship between liabilities and inflation. Intuitively, we would expect, for example, that a rise in the price level of (say) 100% would lead to a rise in the level of rents from real estate and dividends from equities of something like 100% (albeit with complex leads and lags). However, the data do not generally tell us this. The data sometimes even give counter-intuitive answers to some of these economic questions, particularly if the processes have long leads and lags. And the available data are often simply insufficient to allow us to falsify or accept hypotheses relating to the behaviour of the variables over the relevant forecasting period of a generation or two. As Huber has put it, "…it is virtually impossible to evaluate empirically the long-term features of models because of data shortages, which are exacerbated by possible regime shifts or permanent structural changes." (Huber, 1997, page 187).

There has been much debate within the actuarial profession over the types of stochastic models that should be used by actuaries for long-term modelling purposes. It is sufficient to say that constructing a model that is simultaneously parsimonious with the data, has a reasonable accordance with received economic and financial theory and which, if desired, adheres to the efficient markets theory is not easy (and perhaps not possible). Even where longer datasets (for example for UK equities, UK bonds or UK or US cash interest rates) are available, the characteristics of the economic processes may be such that the relevant relationships are not stable. For example, in the last 300 years there have been long periods of stable prices, periods – post war – of accelerating inflation, and periods of stable inflation; when one looks at Treasury Bill return data we find that issues such as whether we include the period of the cash cartel, or wartime – when there were statutory restrictions on interest rates – make substantial differences to the derived model parameters. These problems are not trivial. We cannot simply use the whole data period and say that the average parameters derived from it are okay for long-term use because wars and periods of inflation will occur again in the future. When the data are non-stationary over the long term, average parameters derived from the whole data set do not reflect the behaviour in any relevant epoch. There may be no other way of approaching the problem, but this is, once again, just a warning not to overstate the precision of the results of stochastic modelling.

7.03. Simple Models or Complex Models?

If there is debate about which of stochastic modelling and stress testing is the better technique or whether they are complementary then it is reasonable to debate whether simple or complex stochastic models should be used. This is, indeed, a dilemma that is encountered in practice.

Complex models clearly include more variables and relationships and can, in theory, represent the probability distribution of the financial position of an insurance company more effectively: that is their key advantage. However, given that it is difficult to subject stochastic models for long-term use to the usual processes of scientific falsification, it is difficult to know whether a given complex model is better than an alternative complex model or, indeed, an alternative simple model. Complex models include more relationships – but we cannot easily know whether those extra relationships are either important or valid. However, as a minimum I would expect a model to represent the relationships between all asset classes and between assets and liabilities, paying particular attention to the influence of variables such as inflation that might affect both assets and liabilities.

Models can be complex in the way that they represent economic theory or complex in the way that they are constructed statistically. In my view, there is little benefit from complex models from an economic theory point of view. Many complex relationships used in econometric forecasting models simply do not remain stable over long time periods. On the other hand, many simple and fundamental relationships could probably be incorporated but are sometimes contradicted by the data (e.g. close-to-unit-gain from prices to wages, dividends, rents, equity values and real estate values, and purchasing power parity theory). Often a judgement has to be made as to whether such relationships should be incorporated. In general, I would hope to see a straightforward set of economic relationships which a director, policyholder advocate etc. could question and ask to be changed (as well as the concepts mentioned above, there should also be a straightforward incorporation of efficient markets theory or an alternative).

Complex statistical models may well be able to take advantage of modern techniques to obtain more information from the data. On the one hand, this can be important when trying to construct long-term models from short data runs. On the other hand, the use of more complex statistical techniques can make the models less stable.

With whatever model is used, I think it should be possible to specify its main economic and statistical features (and justify them) in a couple of sides of A4 and explain them in English (as opposed to equations) in a way that is intelligible to the educated non-actuary. If a model cannot be interpreted, there will be problems obtaining general acceptance. Directors will not wish managers to use models that directors do not understand as it will lead to agency problems because of information asymmetries between managers and directors.

A simple model will produce results that are intuitive. To take an example, if the model incorporates purchasing power parity theory over the long term, we would expect the variability of the real sterling returns from a portfolio of diversified foreign equities to be less than the variability of nominal sterling returns. On the one hand, such an intuitive feel allows the output to be more easily understood. On the other hand, the model may miss important features. A simple model will also be more transparent in that it can be seen how different variables interact and lead to a given set of results. Arguably, though, if a model is too simple and produces results that are too intuitive, then we could just use intuition and not bother with the quantitative model! However, this is not strictly true. The model might help us both to quantify and confirm an intuitive hunch (e.g. we may know that greater volatility of long-term interest rates will raise the cost of guaranteed annuity options, but the model may indicate to us the extent of any increase in cost).

Also, it is normally easier to understand the weaknesses of a simple model and quantify the impact of those weaknesses. However, against this, we could argue that if we can understand the weaknesses of a simple model we should perhaps fix them and use a more complex model.

One argument for using a more complex model is that there is an intrinsic complexity in financial relationships. We therefore cannot test one relationship at a time – we can only test one relationship in the context of the wider system which is necessarily complex. Furthermore, if we have not specified that wider system, then tests of a given single hypothesis are not valid. Against this it can be argued that complex models can become over specified and obvious relationships may be missed as a result of a data mining approach that leads to the incorporation of large numbers of relationships that may not be stable over time. A good example here is the Treasury model of the economy that contained 500 equations in 1979 yet had no equation for the impact of real wages on employment or for the impact of money supply on inflation: twelve simple equations would probably have been better.

8.00 Combining Stress Testing and Stochastic Modelling

It is possible to combine together the approaches of stochastic modelling and stress testing, though it is not necessarily easy to interpret the results in terms of probability statements. This is discussed in the recent Bank of England Financial Stability Review referred to above (Bank of England, 2007, page 33). The idea suggested is that the parameters of a value-at-risk model should be changed in certain situations to take account of the fact that at certain times the "normal" parameters may not hold. This should help address the much-discussed problem of non-linearities in times of stressed financial conditions (market movements being more highly correlated than under "normal" conditions). However, as with most stress testing, it is difficult to determine the probabilities of the parameters changing – though in a sense this approach combines both the advantages and the disadvantages of stress testing and stochastic modelling. If this approach were adopted, it might allow simpler underlying stochastic models to be used. Indeed, Robert Clarkson did suggest a simple stochastic model based on a "regime switching" approach (see Booth et al, 2005) to overcome some of the problems that actuarial stochastic models have in representing data patterns realistically. Thus, for example, inflation might be modelled in a particular way and there would be a stochastic input that could change the parameters (e.g. increase the volatility or the mean level of inflation for a period of time). There are some similarities between this approach to stochastic modelling and the approach suggested in the Financial Stability Review (though they are not the same). It certainly has some merits though I am not aware of this approach to model building being at all widely used.

9.00 Conclusion

It is reasonable for actuaries to strive to improve their stochastic models and part of this process will involve increasing the complexity of their models. If I were a director, I would want to know that consistent results were found from both simple and complex stochastic models and from stress testing. None of those approaches can provide confidence intervals that do not have large standard errors (and at times it may be difficult to calculate the standard errors). The models should be used for understanding and, to the degree it is possible, quantifying risk and the distribution of possible future outturns (and, in turn, the amount of capital that must be held to limit the probability of failure) but we should be careful about how we quote the results. Just as there is uncertainty about the outturn of the solvency position of a life office, there is uncertainty attached to our efforts to quantify the uncertainty. This is particularly acute in life insurance company modelling because of the long projection period over which a model is used compared with the period over which the data is collected. Thus it is important to take the results of stochastic modelling together with those from stress testing and more subjective analysis in coming to judgements about issues such as the amount of capital that should be held. We should be careful not to ascribe too much precision to the way in which we quote our results whatever approach or combination of approaches are used. All forms of modelling are part of the attempts to quantify uncertainty and understand the risks facing the business.

References and Bibliography

Alexander C. (1996), The Handbook of Risk Management and Analysis, Wiley, UK.

Bank of England (2006), Financial Stability Report, Issue number 19, Bank of England, London, UK.

Bank of England (2007), Financial Stability Report, Issue number 21, Bank of England, London, UK.

Booth P.M., Chadburn, R., Haberman S., James, D., Khorasanee, Z., Plumb R. H., Rickayzen B. D. (2005), Modern Actuarial Theory and Practice second edition, Chapman & Hall/CRC, London, UK.

Haldane A., Hall S. and Pezzini S., (2006), A New Approach to Assessing Risks to Financial Stability, Financial Stability Paper, number 2, Bank of England, London, UK.

Hayek F. A. (1988), The Fatal Conceit: the errors of socialism, The University of Chicago Press, Chicago, US.

Heffernan S. (1996), Modern Banking in Theory and Practice, Wiley, UK.

Huber P. P. (1997), A Review of Wilkie’s Stochastic Investment Model, British Actuarial Journal, Volume 3, Part 1, No. II, pp 181-210.

Markorov V. (1998), Value-at-Risk and Beyond: the Comprehensive Measurement of Market Risk in Risk Management for Financial Institutions: Advances in Measurement and Control, Risk Publications, UK.

Shaw J., (1997), Beyond VAR and Stress Testing in VAR: Understanding and Applying Value at Risk, Risk Publications, UK.

Shiu Y.-M., Moles P., Adams A. T. and Chan C.-C., (2006), Empirical Evidence on the Use of Dynamic Solvency Testing and Financial Condition Reporting in the United Kingdom Life Insurance Industry, Annals of Actuarial Science, Volume 1, Part II, No. 2 pp 359-392.

Annexe

The main part of this document has focused on the use of stochastic modelling and alternative techniques for assessing risk. This is relevant for various parts of the proposed reattribution transaction. A further issue is the relative merits of stochastic modelling and other techniques for assessing the *expected* value of a balance sheet variable or policy benefit (such as the liabilities of an insurance company, the value of the inherited estate, or expected payouts to policyholders pre and post reattribution in the case of this transaction).

This issue is of relevance to the discussions between the policyholder advocate and Aviva because both parties are using a variety of approaches to assess the value of the benefits to policyholders and the value of benefits forgone from the reattribution of the inherited estate. The purpose of this Annexe is to draw out relevant points regarding the value of stochastic modelling versus deterministic approaches for this purpose.

Many of the points in the main part of this note are still relevant. One should be careful about the degree of accuracy one ascribes to stochastic models. However, because one is using the central part of the distribution of outcomes (normally the expected value), the points made above about the uncertainties surrounding the tails of the probability distribution produced by stochastic models are less important. At the same time, stochastic models may well not give rise to significantly different results from deterministic valuation techniques that are based on expected values of the relevant valuation variables – though this is a generalisation for the reasons explained below.

In general it is not true that the expected value of the function of a variable is equal to the function of the expected value. In other words, if we simply use the expected values of a series of variables and undertake a deterministic valuation it will not give rise to the same result as doing a full stochastic projection of all possible values and using the probabilities of occurrence of different financial outcomes to calculate the expected financial outcome. This is a rather complex mathematical statement that is difficult to put into plain English; however an example illustrates the point. Say we have a die which has three sides with four spots and three sides with six spots. The expected value of the square of the number of spots when we roll the die is 0.5*16+0.5*36 = 24. If, on the other hand, we calculate the expected value of the number of spots first (five spots) and square this we get a different answer (25). Differences between the outcomes of a full stochastic valuation and a deterministic valuation using the expected values of the various variables are likely to be small unless we have options attached to policies. Such options might be explicit or embedded. Embedded options might exist where the company takes the risk that asset shares fall below policy payouts in certain situations. This is, in fact, a feature of the reattribution. Options give rise to non-linear payoffs and the probability distribution of financial outcomes for policyholders is truncated (policyholders have some kind of a guarantee that payouts will not fall below a certain amount – or some other guarantee with a similar effect).

The above point must be taken into account when assessing the outcome of deterministic models, but such models are still useful when valuing payments to policyholders before and after the reattribution or when valuing the inherited estate. Indeed, deterministic techniques might have certain advantages. They often allow one to change variables and see financial outcomes under different assumptions, at a low level of detail, very rapidly. Rather like stress testing, the simplicity of deterministic techniques, combined with modern computing power, brings certain advantages. It should also be possible to reconcile the outcome of deterministic and stochastic valuation approaches. For example, the values of options can be estimated separately when deterministic techniques are used for a valuation and appropriate account can then be taken of them. The assumptions underlying the stochastic and deterministic frameworks can also be compared and it should be possible to ascertain the reasons for any apparently conflicting results using the two approaches.