B.A.J. **10**, I, 153-165 (2004)

ASSET MODELS IN LIFE ASSURANCE: VIEWS FROM THE STOCHASTIC ACCREDITATION WORKING PARTY

A Discussion Meeting

[Held by the Faculty of Actuaries, 17 November 2003]

Mr T. D. Kingston, F.F.A. (Past President): The subject of this meeting is 'Asset Models in Life Assurance', from the Stochastic Accreditation Working Party. Two members of the Working Party are to make a presentation to us, and then the meeting will be open for comments.

Mr C. E. Dowthwaite, F.I.A. (beginning the presentation)**:** I thank the Faculty for the opportunity to speak on what, I think, is a topical area. Many actuaries working in life assurance feel that we are entering a new world at the moment: the world of realistic balance sheets, individual capital assessments and fair value financial reporting. It is absolutely clear that the choice and calibration of stochastic models in this new world are going to be key issues for the profession going forward.

I would not claim that our Working Party has answered all the questions surrounding the use of these models. Indeed, we have not tried to answer them all. What we have aimed to do, instead, is to provide a base from which the profession can work. It is to be hoped that this discussion will be an opportunity to extend our thoughts and the thinking of the profession in this area.

We will be touching very briefly on general insurance and also on life liability models. They are very brief references, but, as these areas fell within the terms of reference of the Working Party, we felt that we should cover them.

By way of background, the Stochastic Accreditation Working Party started as part of the profession's response to International Accounting Standards (IAS), and, in particular, the response to the IAS Draft Statement of Principles for insurance business. As such, it was one of five successors to Chris Hair's Fair Values Working Party (Hairs *et al.*, 2002). Some of you will have been at the presentation, earlier today, of David Hare's paper, produced by another of these groups (Hare *et al.*, 2004).

As well as IAS, we were asked to look at risk-based capital and to cover both life and general insurance liability models. There was a slight change in focus following the IAS decision to split their project into two phases. It was decided to continue with the Working Party's work, as it was clearly relevant for a number of other developments in the life market.

Our terms of reference essentially covered consideration of the suitability of the various stochastic models that are available, both for risk-based capital and for fair value liability assessments. We were also asked to look at criteria for the accreditation of these stochastic models. Within that there was a presumption that there would be an accreditation process of some sort. We felt that it was within our remit to consider whether an accreditation process was appropriate. We were also asked to look at the range of stochastic models that are available, and to have a look at a cost/benefit analysis for the application of these methodologies.

Initially, the major work focus for the Working Party was to have a look at some criteria for the selection of models for use both in fair value and in risk-based capital work. These criteria were to include both the theoretical requirements that any models should satisfy and also the range of practical issues that surround the use of these models.

We felt that an important first step was information gathering on the models available. We felt that there was not really a single data source for the full range of models, including single asset models, multi-asset models, and models both in the public domain and those which are proprietary.

Going back to the accreditation process, we felt that we should consider the role of the Faculty and the Institute in the selection and possible accreditation of models, and, in particular, whether it was appropriate that there should be some sort of accreditation for individual models.

To focus on fair values, initially, the background was the requirement for fair valuation as part of the IAS insurance project. As things stand at the moment, there is a requirement to fair value within IAS 39 which covers the valuation of financial assets and liabilities, and there is likely to be a requirement in phase 2 of the IAS Insurance Standard. The IAS 39 definition of fair value essentially says that a valuation should be market consistent.

Following on from that, the Financial Services Authority (FSA) has indicated that, as far as realistic balance sheets are concerned, it is looking for a market consistent valuation of liabilities. So, there is clearly an overlap between the IAS requirements and the FSA's requirements for realistic balance sheets.

A number of valuation techniques are available. The two main ones are the use of closed form solutions and the use of Monte Carlo simulation. There are two main areas where stochastic models may be useful, or may be required, in fair valuation. The first is where the assessment of liabilities depends on the investment returns or interest rates on a portfolio of assets. The key examples here are with-profits business, where payouts depend upon investment returns; guaranteed annuity options, where the value of the payouts will depend on interest rates; and, finally, unit-linked contracts, where there is a mis-match between income and expenses.

The second area relates to insurance risks. There is a body of theory that says that, if insurance risks are diversifiable, then there is no need to include an extra margin for risk. The IAS, in the initial draft statement of principles, took a different view, and insisted on the inclusion of market value margins. This is an area where the use of stochastic liability models may be a useful tool to assess what level of market value margins are appropriate.

Thinking about model assessment criteria in the context of fair values, the key criterion is simply stated, and that is market consistency, which essentially means the ability to reproduce market prices of traded assets. The key phrase within IAS 39 is that the assumptions used to assess liabilities should be consistent with those used by market participants. One of the key implications is that any model used in this context must be arbitrage-free.

There are two key issues in assessing models for this purpose. The first is the calibration of the model. How is the model calibrated? The other is: to which traded assets is the model calibrated? In particular, what if there are no traded assets with which to test the consistency of the model against the market?

I will come later to one of the key areas that I think exercises people who use these models in practice. This is the choice between models that are based around a risk-neutral approach and models that include deflators.

The use of these models for risk-based capital has a much longer history in the actuarial profession. Companies have been using stochastic modelling to assess probabilities of insolvency, for example, for quite a long time now, so they have a much longer track record. The requirements for the use of these models have moved on recently as a result of the FSA's requirements for Individual Capital Assessments (ICAs), which are outlined in CP195. The preference, in the consultation paper, for the use of capital models means that companies have been pointed towards the use of these stochastic models in the context of ICA. A very similar approach is also implied for general insurance. For ICA, the key modelling requirement is to assess capital requirements arising from a number of different risks, with the key risks being market risk, credit risk and insurance risk, in the language of ICA and CP195.

Taking market risk first, asset models may be needed to look at the market risk associated with holding assets such as equity and property, and also interest rate risk, associated with fixed-interest assets. Use of these models is appropriate in any line of business where there is a material asset/liability mis-match.

In the area of credit risk, we look at credit standing and defaults for corporate bond portfolios. Here, the key business areas are clearly annuity portfolios, when the company uses corporate bonds to back annuity liabilities; general insurance, when the insurer holds a lot of corporate bonds; and with-profits business, when with-profits funds are largely backed by fixed-interest portfolios.

Insurance risk is the final key area for risk-based capital. Here we are looking to model the incidence and amount of insurance claims. The key business areas are general insurance and life business when mortality is significant, so that both protection business and annuities are important areas.

I now consider assessment criteria for risk-based capital. The Working Party's view here was that there were no simple tests for model suitability. In particular, in contrast with fair values, market consistency is not required. The distribution of outcomes for the model should be as realistic as possible. That means that risk-neutral models, which are commonly used in fair value assessment, are not appropriate for use in risk-based capital work. The key focus, when assessing whether a model is appropriate for risk-based capital, is usually the tail of the distribution. Again, this contrasts with fair values, where the mean of the distribution may be more important. It is also important to look at the correlation between different risks. We have a number of different risks potentially included in a risk-based capital assessment. Clearly, there are correlations between market risk, credit risk and, in some cases, correlations with insurance risks as well.

Considering calibration; again there is no simple question to answer here, in the sense that there is not a simple criterion for the calibration in the way that there is for fair value-based models. Here, our view was that any calibration should be based on a combination of factors: historic experience; market data; but also that there should be a significant element of judgement involved in calibrating a model for risk-based capital. One of the key issues in calibration is the choice of the calibration period, and, in particular, what period of history to examine. There is obviously a range of potential timeframes to consider when calibrating a model for this purpose. There is also the issue of the periodicity of return to which to calibrate.

On risk-neutral versus deflators, for those not familiar with the concepts, essentially, in a risk-neutral model, the assumption is that all assets earn the risk-free rate on average, and, consistent with this, the discount rate used to assess liabilities is the risk-free rate. In essence, a risk-neutral valuation is a mathematical tool to put a fair value on an asset or a liability.

In contrast with this, in a model which includes deflators, the investment return models are, typically, more realistic. In particular, they could include, for example, equity risk premia, but then the valuation of those returns and cash flows is performed using deflators which incorporate, not only discount factors for the time value of money, but also a valuation of the riskiness that has been involved in earning that return. So, the deflators are used to put a fair value on these real-world cash flows.

The conclusion of the Working Party was that either approach could give acceptable answers for a fair value assessment. The two approaches are mathematically equivalent. I think that it is true that risk-neutral models are more commonly used in valuations. The valuation of options, for example by investment banks, typically uses risk-neutral models. There is an element of convenience in using a model which includes deflators. We have a set of real-world scenarios which can be used for two purposes, for risk-based capital and also for a fair valuation. In essence, our view was that the choice was one of convenience and of practicality, rather than one of theory.

References

Hairs, C.J., Belsham, D.J., Bryson, N.M., George, C.M., Hare, D.J.P., Smith, D.A. & Thompson, S. (2002). Fair valuation of liabilities. *British Actuarial Journal*, **8**, 203-340.

Hare, D.J.P., Craske, G., Crispin, J.R., Desai, A.J., Dullaway, D.W., Earnshaw, M.A., Frankland, R., Jordan, G.A., Kerr, M.G., Manley, G.M., McKenzie, J.L., Rae, R.A. & Saker, M.C. (2004). The realistic reporting of with-profits business. *British Actuarial Journal*, **10**, 223-293.

https://doi.org/10.1017/S135732170000458X Published online by Cambridge University Press

Mr M. Abbink, F.I.A. (continuing the presentation)**:** I shall talk about a number of the models that the Working Party considered. The intention is not to give you a full technical presentation on the models, but I think that it is useful to go through a few of them. More details will follow in the paper.

The area in which there has been the most academic research is that of equity models. The benchmark is still the Black-Scholes model, with which we are all familiar. It has existed for more than 25 years, and it is a very convenient model, because it provides closed form solutions for option prices. However, its disadvantage is that it does not give real-world projections of equity returns. In particular, the tails are not covered very well in the assumed lognormal distribution.

A number of other models have been developed around the Black-Scholes model:

— One of the developments has been, for example, to take the drift (or the mean return) of the equity model as stochastic.
— Another development has been to take the volatility of the Black-Scholes model as stochastic, so that we get fatter tails.
— Another type of equity model which has been used in a number of models is regime switching. This is also sometimes used in multi-asset class models, which Mr Dowthwaite will talk about later. The idea is that equity returns are modelled from two different types of distributions: one distribution is benign, with quite high equity returns and low volatility; and the second is a distribution in a crisis scenario, with low equity returns and high volatility. There is a transition matrix between the two different distributions. It is a very clever way of trying to make the distribution of equity returns more realistic.
— Another type of model, which we have come across when looking at multi-asset class models, is the jump diffusion model. That is a particular type of model where, on top of the normal distribution, you put a Poisson type of distribution — which generates a crash, or jump, every now and then.
— The last type of model that we came across were the GARCH models.

Our intention, apart from giving an overview, is also to look at some of the desirable characteristics of these models in the context of either fair values or realistic balance sheets. Obviously, one important characteristic is that the asset models are able to replicate the prices of options. In that respect, it is very important that the model is easy to calibrate and is transparent. A closed form solution is another desirable characteristic. Obviously that is difficult, because the further you move away from a standard Black-Scholes model, the more difficult it will be to have a closed form solution for your options. To have fat tails is a further characteristic to be desired. This is usually achieved by changing the volatility structure. Something which will become increasingly important for equity models in the context of realistic balance sheets, for example, is that the equity model is able to calibrate to the whole term structure of volatilities.

What we see in the market is that options with different strike rates have different types of implied volatilities. This simply reflects the different prices which market traders ask for these types of options for different strikes. Sometimes it appears that the wrong parameter and implied volatility are put into the wrong model, a Black-Scholes model, to obtain the right market price. In any case term structure volatility is important, because, if you are trying to price an out-of-the-money option, it is important that you calibrate to an out-of-the-money volatility.

Dividends are potentially important, but have been more or less ignored in academic research, which has been focused on modelling total returns for equities; but dividends have a different tax base in life assurance and generate cash flows. So, in the stochastic models that you are building for realistic balance sheets, it might be very important to have a realistic projection of dividends.

In addition to equity models, there are interest rate models. Again, a wide range of models is available. One of the first models developed was the Black's model, which is more or less an extension of the Black-Scholes model. The idea is to price a bond option at a particular point in

https://doi.org/10.1017/S135732170000458X Published online by Cambridge University Press

the future. You simply assume that the interest rates are lognormally distributed at that point in time. The model does not give you any information about how interest rates move between this point in time and the point where the option can be exercised. So, more sophisticated models have been developed over time, particularly the so-called short rate models.

Rather than simply looking at the cash rate, short rate models have an implied term structure. There are two types: one is the Vasicek model, which was one of the first ones; and the other is the Cox-Ingersoll-Ross model, which is an interest rate model which could prevent interest rates from going negative. That is an area which could be important when pricing interest rate options with a very low floor. These two short rate models, although they generate a term structure, are not able to calibrate to that term structure, which is very important, given that a small error in the modelling of the term structure can give you a big error in the option price generated.

The Hull and White model, basically, is an extension of the Vasicek model, where the mean reversion level has become more sophisticated to allow calibration to the whole term structure. Even that model is not sophisticated enough for some purposes. In particular, some of the real world distributions of interest rates could not be captured, and people have now moved onto multifactor models for interest rates, where some of the mean reversion levels of interest rates, for example, will become stochastic in themselves.

All of these models can be captured within one framework: the Heath-Jarrow-Morton framework. In that framework you calibrate the whole initial term structure, and the different models, mentioned above, can be applied by setting different types of functions for the volatility structure of the different interest rates on the term structure. Heath-Jarrow-Morton models more or less cover the whole range, and these are more or less the standard models in interest rate modelling.

There are several desirable characteristics for interest rate models:

— One area which is important for realistic balance sheets is the choice of yield curve. It does make a difference to the numbers. What you are looking for is a risk-free rate. In the option markets people look for liquidity. In the investment bank world people often take the swap curve as being a benchmark for risk-free rates. The International Accounting Standards Board (IASB), in its draft statement of principles, takes government bonds in the first instance, and adds a small margin, but there is a range between government bonds and swaps, and it is important that some consistency is achieved.

— Quality of fit to the yield term is important in interest rate models. Using the Heath-Jarrow-Morton framework, models which can capture the whole term structure obviously have to be preferred to models which can only model, say, a ten-year yield and a short rate.

— Volatility structure is also important. The stochastics of the model and the option price that you derive depend on the volatility term structure that you assume.

— It is also important to look at how the different yields on the yield curve correspond with each other. Some of the models, for example, are not able to generate all the different shapes of the yield curve.

— Then there is ease of calibration. As with equity models, it often leads to faster calibration if you are able to price your options with a closed-form solution.

The last area that I will cover is that of miscellaneous asset classes. We looked at some property models, although, unfortunately, there is little on property models in the literature. There are some econometric models, but it seems to us that people, when pricing property risk, will normally use an equity model and change the parameters. Some calibrations are available, but it is very difficult to get a market-consistent calibration, and most of the time it is based on historical data.

We next considered corporate bonds, which is an extension of the consideration of interest rate models. Although it is a very important area, model development has been quite slow. Most of the models that have been developed by banks in the last two decades have focused on pricing one particular bond. The way in which banks do that is to set up internal rating classes and try to predict a default for each bond individually. That is not necessarily a market-consistent approach, and is not necessarily ideal for pricing. We looked at some models, again extensions to the Black-Scholes approach, that will give you an option price approach for corporate bond risk.

Overseas assets is another category where we tried to find some literature. Again, that was very difficult, and is an area for further development. The key question is whether overseas equities have a different distribution from domestic equities, or whether you can use the same distribution, but, perhaps, with different parameters.

The last area which is worth noting is the complexity of portfolio options. To do a realistic balance sheet calculation correctly, you need to take into account, not only the proper characteristics of your liabilities, but also the options that you own in your assets. There is a conflict, because it is difficult enough to price all the different options in your liabilities correctly, but if you also need to price the options in the assets that you hold, for example to cover guaranteed annuity options, then it becomes even more difficult.

Mr Dowthwaite (concluding the presentation)**:** Although the Working Party's main focus was on asset modelling, it was within our remit to have a look at liability models. On the general insurance side, our conclusion was that there is a wide range of models available commercially. Because of the range of models that are available, we felt that it was not appropriate to investigate those further. There is clearly a range of quite complex models that are widely used.

In contrast, on the life insurance side, we felt that there was a relatively limited number of mortality models available. There are some demographic models. Lee-Carter is a prominent example of a demographic model that is used to project mortality forward. We felt that this was an area where, although there is quite a lot of activity at the moment, there was potential for further work and further development within the profession.

One key issue that we saw here was the issue of consistency with deterministic mortality projections, and, in particular, consistency with the CMI scenarios. If we are using these models to develop risk-based capital for life assurance portfolios, then it was important that we understand what the relationships between the CMI projections are and the output from these models.

We felt that there was also a potential need for some work on morbidity models and also for models that linked persistency to economic scenarios. Again, although there has been some published work on dynamic lapse rates, this could be an area for further development.

To move on to multi-asset models — for many of the practical exercises in which actuaries become involved with asset models, and in particular with realistic balance sheets, the use of a multi-asset model is essential. That is for any situation where liabilities depend on the performance of a portfolio of investments, particularly the assessment of with-profits liabilities. In many cases these multi-asset models build on the individual asset class models about which Mr Abbink has talked, and so consideration of those individual asset class models is quite important in any assessment of a multi-asset model. A multi-asset model, though, clearly needs to add in correlation between individual asset classes. One particular issue here for fair value work is that there is no market price for that correlation in most circumstances. It may be necessary to use historic correlations as a guide to correlations within the model, even within a fair value context.

It became quite clear to us that there is quite a wide range of both public domain and proprietary multi-asset models available, and we felt that one of the things that we needed to do was some information gathering on these models. As a number are proprietary, we felt that it was appropriate to run a survey. We sent out a fairly detailed questionnaire to all of the owners of the models of which we were aware. I thank those who completed the survey. It was quite a long list of questions, and I appreciate that it was quite an extensive piece of work for those who participated. We had detailed responses on 12 models. The questions covered quite a wide range of topics, including the range of assets modelled, the model's ability to reproduce the market price of assets, hence the suitability of the model for fair value work, and some questions about the tails of distributions within the model, so looking at the models from a risk-based capital perspective.

We also asked a number of practical questions about the model: "How long has the model been in development?" "Are there any practical issues about which users of the model need to be aware?"

It became clear that this is an area where there is quite a lot of work going on at the moment. In a number of cases there were models that were in the process of being quite significantly extended and developed. We felt that it was not appropriate to publish details of the survey within, for example, a Working Party paper. We came up with the idea of making the survey results available, possibly through the Faculty and Institute website, so that people working in this area have a relatively up-to-date single source of information on the wide range of models which are available. It is important, though, that there is an update process for such information, because models are being developed all the time.

To move on to some of the practical issues, for most exercises, particularly those involving multi-asset models and those involving with-profits funds, run times remain a key issue. The number of scenarios required to get an accurate assessment, for example of a realistic balance sheet, depends on the model chosen, and there are key differences between the number of scenarios required for different model types. There is a wide range of variance reduction techniques available, however, which enable the user to reduce the number of scenarios required.

It is also worth commenting that there are some deflator-based models that can be specially calibrated to try to give a reduction in the number of scenarios required.

Calibration is an absolutely key area. The availability of assets to support the fair value calibration is one vitally important issue so far as calibration is concerned. In particular, long equity options are not widely available in the market. There is not a deep and liquid market in such options. So, the assessment of long-term equity volatility and the parameters used for long-term equity volatility are a key issue, for example, in realistic balance sheet assessments. Also, the ease of calibration is a key issue. Some models may be easier to calibrate than others.

We felt that one final practical issue was understanding the model. This is a relatively new area for many actuaries. Some of the models that are in use, particularly the multi-asset models, are quite complex, and we felt that it was important that anybody using these models in practice should understand the uses of individual models and also any limitations that particular models have. We felt that it was important that a model should be as understandable and transparent as possible for potential users.

To come to the stochastic accreditation process, there was a presumption in the terms of reference of the Working Party that there would be an accreditation process of some sort, and that there would be a Stochastic Accreditation Board working under the auspices of the Faculty and the Institute, which would give some sort of kitemark to particular models. Our thoughts on this were, first of all, that the suitability of the model must depend on the application. In particular, the choice of the model must be made in relation to the liability being valued.

To give an example, some of the simpler equity models that we talked about earlier are appropriate for the valuation of equity options. However, these models would not be appropriate in the context of a full risk-based capital assessment, or a realistic balance sheet for a with-profits fund that includes a wide range of assets. Clearly, appropriateness has to be measured in relation to the liabilities being assessed.

We also felt that calibration was absolutely vital. It is clear, for example, in fair values, that the calibration of a model must be such as to reproduce market prices. So, accreditation of a particular model is not sufficient, because calibration has to be appropriate, too.

We felt also that an Accreditation Board could operate as a disincentive to innovate. We felt that it could result in an element of inertia, and a tendency to stay with existing accredited models, which could potentially deter those who were looking to innovate, to introduce new models, or build on existing models.

We felt that there were a number of practical issues surrounding the operation of a board.

The first one was that there could be a time lag between models becoming available and models being accredited. This is particularly an issue in the context of the speed of change in this area. There is a lot of development going on with these models at the moment, and anything which slows down the process could be a significant drawback.

We felt that there was an issue in funding the operation of the board. It was not clear who would fund such a board or, indeed, who would actually provide the membership of such a board.

In the context of the use of these models, for example in fair values, and particularly in the context of fair value accounts, we felt that the final sign off of any accounts would rest with the auditors, and it was not clear to us how much reliance auditors could place on the fact that a particular model had been accredited by a board sponsored by the profession.

To cover the Working Party's recommendations, first of all, as you have probably guessed, we believe that accreditation of individual models is not appropriate. We did not feel that, given the issues which we have just gone through, it was appropriate that the Institute and the Faculty should set up an Accreditation Board to provide some form of accreditation to individual models. We did feel, however, that professional guidance for actuaries working in this area would be helpful, and, in particular, that guidance would be helpful in the areas of choice of model and also in calibration. Guidance also could be helpful when looking at some of the practical issues which actuaries who are working in this area face. One example would be looking at the number of scenarios that are required to derive an appropriate assessment of liabilities.

Regarding the recommendation that there should be a successor working party to develop this guidance, I am pleased to say that this is likely to happen under the auspices of the Supervision Committee. I believe that a further working party will be being set up in the very near future.

One final idea from the Working Party was that a central data source could be useful in model calibration. We felt that quite a lot of data is required to calibrate models, and having a single Institute and Faculty sponsored data source for that information could be quite a useful step forward.

So, to round off, I said in opening that we have certainly not tried to answer all the questions surrounding the use of these models. We hope, however, that we have provided a platform for the profession to work from in thinking about some of the issues.

Professor A. D. Wilkie, C.B.E., F.F.A., F.I.A.: The presenters made a distinction between what I would call *real world models*, which are necessary for risk-based capital, and *option pricing models*, which are appropriate for calculating fair values and, of course, option prices. Real world models can be made as complicated as you like. You can put in jumps or fat-tailed distributions, and also stochastic parameters, which I would call a hypermodel, choosing the parameters for each simulation randomly each time that you run it from some parameter distribution; different scenarios, jumping from one to another; regime switching.

On the other hand, with an option pricing model, if you are using it for fair values, the important thing is to obtain the hedging quantities that you need in order to hedge that particular option. The option price is the right price for the option, because, and only because, it is assumed that you can do the hedging process. Therefore, I am very much in favour of using closed form solutions for option pricing rather than deriving them from stochastic simulation. Stochastic simulation, whether with real world models or deflator models, does not so easily seem to give you the hedging proportions.

The way that I would want to look at this is the way that was described in the paper Wilkie *et al.* (2003) on guaranteed annuities, and presented here in January 2003, to work out an option pricing formula using some option pricing model, and then embed that in a real world model to see what your hedging error might be, see how well it matches up, and then base your risk-based capital quantile reserves, the contingency reserves, on the hedging errors. You can also bring in the uncertainty that there is about the parameters that you are going to use. Whatever option pricing model you use, you do not know whether those parameters are going to apply throughout the period of the option, particularly over long periods. So, I favour using complicated real world models and simple option pricing models, especially ones that allow closed form solutions. That is from the point of view of a valuation within a life office. That is a different objective from that of traders of options, who are setting prices, buying or selling options, and who may want a more elaborate option pricing model. They also ought to be concerned with the risk-based capital that they may need, but their primary purpose is to sell the individual traded options, whereas, within an insurance context, the options are often embedded in a much larger type of insurance policy. It seems to me that an error in the pricing of them may be small in relation to the total cost of the policy. So, one needs to put it in context.

No models that I have ever seen replicate the prices of all traded assets. I have never seen a model, for example, that can replicate the price of individual shares or of individual corporate bonds. Nobody is attempting to write that type of model, at least not a stochastic model. One could probably get closer with bonds than with individual shares.

One of the ideals of a model — and, of course, it is desirable if it is an option pricing theoretical model — is that it should be arbitrage-free. It is not always clear that the market itself is always arbitrage-free. There may be occasions, particularly because of trading limitations, not being allowed to go short, or it being too expensive to go short, when the market itself may not be arbitrage-free. Indeed, I have been told (and I do not know whether this is still true) that there are occasions when it may be cheaper to buy an option on an individual share and exercise it immediately than to buy the share. That anomaly may last for 15 minutes or so, but it is worthwhile taking advantage of it if you can spot it.

A risk-based capital model could, and I think probably should for this purpose, start from market values or market conditions, and therefore it does necessarily replicate the present market values of the assets that it is modelling. A complicated real world model is not intended to be an option pricing model, so it does not necessarily replicate all the derivative securities that you can get. It obviously can replicate the values of the basic options in which it is trading. The way of trying to combine these two is to use a yearly, or monthly, discrete real world model, and join them together over the periods with stochastic bridges, which again we described in Wilkie *et al.* (2003).

A small technical point is that I think that the risk-neutral model does not necessarily assume that all assets earn the risk-free rate on average. I think that it assumes that all assets earn the same rate as the asset that you have chosen as the numeraire. You can choose any asset as a numeraire, and everything else is a martingale with respect to that numeraire. You could choose shares as the numeraire, and measure everything in share price terms, in which case shares are risk-free, because you are measuring in numbers of shares. If you own 100 ICI shares, you own 100 ICI shares, regardless of how much money that will give you. So, it is money that changes in terms of ICI shares and not the number of shares that you own. You can choose any numeraire.

I am ignorant about how life offices have developed model points to represent the portfolios, but a way of economising in computer run time would be to have a good structure of model points to represent the portfolios, which cuts down the number of cases to which you are applying your model. I am sure that it is something that is done widely, but not something that has come forward in theoretical papers very much since Lidstone presented the *n*-ages method about a hundred years ago, which was simply a model using model points.

I think that a central data source is very good, and I advertise the fact that Andrew Cairns and I have developed a gilts database which is on the Heriot-Watt University website. It gives the value of indices at monthly intervals since the current series of *Financial Times*-Actuaries bond indices started, along with corresponding gilts prices. I think that it is rather sad that the FTA and FTSE Actuaries indices, which have been sponsored by the profession for the last 41 years, have not actually been well recorded. The Faculty and the Institute do not have historic records. The *Financial Times* used to produce a big black book that gave prices twice a month. It was full of misprints! There is plenty of information that we sponsor through the FT indices, and we actually do not keep it. A good data source would be extremely useful. The one that I have mentioned is just gilt prices. One would like to see swaptions and option prices recorded as well, to the extent that one can.

Mr W. P. McCrossan, Hon.F.F.A.: I will speak only about overseas markets. In Canada and the United States of America we have been insuring the stock market for about 35 years, and have some good and bad experiences.

I will explain how and why we entered the market, and then come to the issue of stochastic models. We entered the market because the Securities Commission said that life insurers could not sell unit-linked products unless they offered guaranteed death benefits and guaranteed maturity benefits of at least 75% of the money paid in. So, since the life companies wanted to sell these products, they entered the market, starting in 1968. The question of how to price death benefits and maturity benefit guarantees was the subject of a landmark paper in the *Transactions of the Society of Actuaries*, written by Frank De Paulo, who was then the Actuary of Confederation Life. He had a theory, which has worked out reasonably well, that, by looking at the probability distribution functions of the changes in the stock market, one can capture human behaviour and emotion; that is, markets are driven by fear and greed, not necessarily by economic fundamentals. He mined all of the data that were available on the Standard & Poor's and Toronto Stock Exchange indices from the late 1920s to the then present, which was the late 1960s. The products turned out to be reasonably profitable through that period. Then life started to become interesting about ten years ago, because insurers had made money on the products, and they started to improve the products. They started to move from 75% guarantees to 100% guarantees. Then they started to move to guarantees that were re-settable, so that, if you had a certain gain in the market, you could re-set the guarantee period and start over again.

Then, some really interesting divergence of opinions started to appear in the marketplace. The divergence of opinion had to do with market validation, if you will, and the choice of relevant data. Essentially, two camps were developed. I was in the camp where all the data that are available are regarded as relevant, because human behaviour has not changed very much. Therefore, one should look at the probability distribution functions from the late 1920s, an approach which had also had the advantage of having two market up cycles and two market down cycles.

The other camp decided that only post-war data were relevant. They rationalised that the old economy (the depression, and so on) and the time prior to the existence of the Securities and Exchange Commission were really irrelevant to a modern economy. In effect, there was a paradigm shift, just like the dotcom shift of the 1990s.

What became interesting, as people used these two paradigms to price products, was that there was a very big discontinuity between those who used all of the data and those who used post-war data, because the tails are much fatter when you use all of the data.

Around about 1994, the tails became so fat and the differences became so big, that what happened was that it became pretty obvious that the reinsurers were using post-war data and the direct writers were using all of the data available. So, the direct writers were saying: "This death benefit should cost me 80 or 100 basis points, but the reinsurers will take it off my hands for 40 or 50." So, they very quickly worked out that the optimum strategy was to offer the highest benefit for which the reinsurers would insure 100% of the risk. This, effectively, became the marketing strategy of the direct writers from 1995 to 2000. In 1999, or thereabouts, the reinsurers started to pull back from the market, because they realised that they were taking a significant systemic risk.

I presented a paper on this topic at the joint meeting with the Australian Institute of Actuaries in August 2003, and it has also been the subject of a two-day seminar at the Canadian Institute of Actuaries. Using data as far back as possible, that is the 1920s, we accurately picked the boom of the 1990s and accurately predicted the fall of the early 2000s.

When you talk about market validation, it is very interesting to look at the difference between the prices of contracts using the two different sets of probability distribution functions: those used in the pricing; and those which have happened subsequently. Treating the changes in equity price as a function of emotion, as opposed to a function of economic logic, seems to track better what actually happens in practice.

Mr J. A. Jenkins, F.I.A.: I agree that accreditation is not appropriate, and would not work. If it were to work, we would need an independent team of people who were separately remunerated, which I do not think will happen. If we just do it in our normal way of people doing it in their spare time, the proprietary nature of some of these models and conflicts of interest will get in the way. I just do not think that that system will work.

I think that there is an alternative. We need guidance setting down minimum standards for the characteristics necessary: calibration; to what; how closely; how to take into account the specific circumstances of the company; what volatility, as a minimum, to ensure that the models are not too kind; and what sort of disclosure is needed for the approaches used. I think that such guidance could be provided either by the profession or by the FSA. I would prefer the profession.

If you look on their website, you will see that the South African profession has taken a lead. It already has a draft guidance note on stochastic modelling, which, I believe, does lay down some characteristics and minimum volatilities, etc.

I have a supplementary question to those of Professor Wilkie. I note his comments on the closed form solution for valuing options and guarantees under fair values. What I do not understand is how he would allow for management actions (for example changes in investment mix) to reduce the likelihood that the guarantee will actually apply.

Professor A. J. G. Cairns, F.F.A.: The first matter that I should like to discuss briefly is the issue of risk-neutral models versus deflator-based models. There have been many debates and misunderstandings over these two approaches. It was certainly good to see in the presentation the statement that both approaches are mathematically equivalent. For example, if you start off with a model that is based in a risk-neutral framework, then you can always modify that model to re-express it as a deflator-based model, and vice-versa, so you can always start off with a deflator-based model and convert that into a risk-neutral model. Sometimes that is very easily done, but there are many other cases where a conversion from one framework to the other is not so easy. So, when you are building up your model, it depends on which sort of framework you prefer.

Life office actuaries who are thinking about buying one of these models should certainly not be thinking: "I want a deflator model", or: "I want a risk-neutral model", because both types of models are mathematically equivalent, and will give the same answer. So, what you must be looking for is which model is the best model that is available for the particular task that you have in hand, be it fair value of long-dated liabilities on the one hand or a short-term risk-based capital type of approach on the other.

Something was said earlier which I thought was perhaps slightly misleading with regard to the risk-neutral type of model, and that is that risk-neutral models, *per se*, are not appropriate for real-world calculations. Of course, that is certainly true if you are doing your simulations using risk-neutral probabilities, but, in fact, if you are using a risk-neutral model it is very easy — and risk-neutral models have the facility to bolt an extra layer on to that model, where you add in market prices of risk and risk premiums on different assets, so that you can re-run your model with real-world probabilities and appropriate risk premiums. You can also build in mean reversion, if that is what you want.

The next point that I want to make is about understanding your model. This is a very important point for people who are going to be using these sorts of models. You must not apply these models blindly. You have to understand what is going on in the model, and which aspects of the asset model are important for the liabilities at which you are looking. For example, I have seen it suggested that, in certain packages that you can buy, all you will get is 10,000 scenarios given to you on a CD. To my mind, that is totally inappropriate for the work that you are doing. You cannot possibly understand what is important in the model that you are using. What you really need to be able to do are such things as sensitivity analyses. These procedures are only possible if you have the model itself on your computer, and not only its output. This allows you to change the parameters that go into the model, and you can understand what it is in the model that really influences the results of your liability calculations, or the calculations of quantile reserves if you are looking at risk-based reserving.

Considering long-term risk versus short-term risk, there was mention earlier of the issue of fat tails, for example. Again, this is an issue of which is the right model for the particular application at which you are looking. If you are looking at a short-term assessment of risk, I think that having fat tails in your model is very appropriate. On the other hand, if you are looking at a longer-term problem, such as fair value, then fat tails are much less important as a component in your model.

What, perhaps, is much more important in the longer term are the issues of model and parameter risk. Professor Wilkie mentioned hypermodels, where, at the start of each simulation run, you re-simulate the parameters that you put into the model. In my mind, that is certainly much more important, particularly parameter risk, than whether or not you have fat tails in your model.

My final comment is with regard to a database which could be made available to everybody. I think that we could add to that the issue of how actuaries should calibrate their models to get market consistency when we are looking at longer-term options. Certainly, we can go onto the LIFFE website and get prices for implied volatilities for short-term options, but, on the other hand, where do we get these numbers if we are looking at a 20 or 30-year equity option? If there is a central committee which could look at that question and find appropriate sources of data, it would be very much appreciated.

Mr T. D. Kingston, F.F.A. (Past President): I approach this subject as somebody who has rarely used models. I am now a non-executive director or an adviser to a wide range of financial institutions, including reinsurance companies, life companies, non-life companies and pension funds. Across all of these we are seeing increasing use of asset/liability models. I guess, initially, that most of these were liability models. Where I saw a lot of them, to begin with, was on the reinsurance side, particularly in some of the extreme risks in non-life. On the general insurance side, we are certainly now seeing a lot more interest in asset modelling and carrying that through to asset/liability modelling. We have heard a great deal in this discussion about life insurance. I will not add anything to that. It is on the asset side where we are probably seeing the most sophisticated modelling.

Something that has disturbed me — I have become reinvolved with pension funds in recent years, having not been involved for 20 or 30 years — is how far behind the pensions industry is in terms of asset/liability modelling, and, indeed, in the whole structure of risk. We have seen the consequences of that in the last two or three years, where pension funds have swung from a very healthy situation to a very unhealthy situation. The actuarial profession does not come out of that very well in terms of its failure to predict at least the possibility of that happening. So, I hope that we will see more about that in future.

Perhaps one of the things that emerges is the whole question of sharing experience across the different aspects of our profession. What we have had in this presentation and discussion has come out of the life insurance business. I hope that we will be able to share that across other aspects of our business, because it is rather surprising that there has been so little sharing.

The Finance and Investment Board is sponsoring a survey of what people are doing in practice in asset/ liability modelling across a range of different financial institutions. That work should be starting relatively soon. I hope that its results will come out in 2004. I hope that it will be able to combine with the successor to the Working Party which has been reporting to us here, and that we will get more from that as well, because there is no doubt that this is an important and growing part of our work, and it is not something that is going to go away.

Mr C.E. Dowthwaite, F.I.A. (replying)**:** Concerning Professor Wilkie's comments on closed form solutions, I agree entirely that, if we can use closed form solutions, then that is a desirable option.

I share some of the concerns of Mr Jenkins over management action. One practical issue that life office actuaries will be facing in the very near future is projecting realistic balance sheets and the issue of doing stochastic projections within stochastic projections. I suspect that, for the foreseeable future, life office actuaries will have to rely on closed form approximations when projecting realistic balance sheets.

I found Mr McCrossan's comments on market behaviour very interesting, and I shall certainly look out his paper.

Then, in response to Professor Cairns' comments on the use of risk neutral models for real world calculations, I have to admit that we have been slightly inexact in referring to risk neutral models — we should, of course, have referred to models with risk neutral calibrations.