ABSTRACT OF THE DISCUSSION

HELD BY THE INSTITUTE OF ACTUARIES

Professor A. D. Wilkie, C.B.E., F.F.A., F.I.A. (introducing the paper)**:** The origins of this paper lie in the mid-1960s when Sidney Benjamin wrote a number of papers (Benjamin, 1963, 1966), on how actuaries could use computers, and included the concept of simulating a range of possible futures in order to calculate contingency reserves for various types of life assurance. These ideas were developed further in the 1970s by the work leading up to the Report of the Maturity Guarantees Working Party (MGWP) (Ford *et al.*, 1980), to which Sidney Benjamin and I both contributed. The stochastic model for shares developed for the MGWP stimulated me to produce the first version of my stochastic asset model in 1984.

In January 1985, I presented a paper to the then Institute of Actuaries Students' Society (Wilkie, 1986). In it I used a number of examples to show how my stochastic asset model (or any other similar model) could be used to derive interesting results. I did not use guaranteed annuity options (GAOs), though they would have been an excellent example for my purpose. If I had, Section 2 of the current paper could have been written at that time. I do not know why I did not consider GAOs at that time. I suppose that they were not of topical interest then. In fact, the life office for which I then worked had issued rather few such contracts, and quite a long time previously.

It was not until about 1997 that the topic surfaced publicly. Bolton *et al.* (1997) produced a useful, and too little publicised, short report on the subject, that contained all the essential principles. Then Sheauwen Yang arrived at Heriot-Watt University, looking for a subject in which to do a PhD, and GAOs were immediately to hand. She completed her thesis in 2001, and this paper is derived from her work, though rewritten for an actuarial audience. Her thesis and our paper are not, however, the same. She covered aspects such as monthly premium policies, with-profits policies and stochastic mortality, which we do not discuss fully, and we cover the history since 1985 more fully, and use a more elaborate option pricing formula than she did.

The paper was presented to the Faculty in January 2003. If it had been written solely for this meeting, we would have shown examples for the end of 2002 for entrants in 2003. However, since the long-term interest rate (in fact the yield on the FTSE-Actuaries irredeemables index) was, at 4.56%, almost the same as it was at the end of 1998, the results are similar, if one allows for some improvement in mortality since then.

We might also have used the CMI's latest mortality forecasts, based on the cohort effect. However, all that these would show is that, if mortality improves further, GAOs at any fixed rate are even more 'in the money', so more expensive, whether one uses the static hedging or the dynamic hedging approach to investing the assets representing the reserves.

We include several concepts in this paper that we consider to be of general application in considering reserving for the liabilities of any financial institution, not only a life office. There are three quantities of interest: best estimate, prudent reserve (including a contingency reserve), and fair value. Fair value lies between the other two, and allows for some return on the capital that needs to be provided to set up the prudent reserve. Premiums should resemble fair values. They are almost the only market prices that we have.

There are two approaches to the investment of the relevant assets: a static investment policy, as we use in the first sections of this paper; and a dynamic investment policy, that is dynamic hedging, based on an option pricing methodology. Whichever one uses, prudent reserves also need to be calculated, with fair values based on them. One cannot assume that the real world will behave exactly like any theoretical option pricing model, any more than one can assume that it will behave exactly like any traditional deterministic actuarial model. If those who read the paper take at least that lesson away, we shall feel that our work has been worthwhile.

References

Benjamin, S. (1963). Computers and actuarial science. *Journal of the Institute of Actuaries Students Society*, **17**, 93-153.

Benjamin, S. (1966). Putting computers on to actuarial work. *Journal of the Institute of Actuaries*, **92**, 134-192.

Wilkie, A. D. (1986). Some applications of stochastic investment models. *Journal of the Institute of Actuaries Students Society*, **29**, 25-51.

Mr D. W. Dullaway, F.I.A. (opening the discussion)**:** The authors have produced a very interesting paper. In considering how to reserve, price and hedge guaranteed annuity options (GAOs), they bring out, in microcosm, almost all of the important issues that life offices face today, at least from a technical perspective.

In the paper the authors have adopted a historical approach to the GAO problem. They discuss how an actuary might have addressed this issue in 1985, show how well (or badly) this approach would have worked since then, and finally ask how we might address a similar problem today. This gives an interesting picture of how mainstream actuarial thinking has developed over the last 20 years.

I do not wish to dwell on the past, but to consider what this paper tells us about the future, and how we should now deal with GAOs and similar options today. The starting point is Section 10, where the authors introduce option pricing and hedging. Considering pricing first, and ignoring mortality risk, if a market price for an appropriate hedge exists, then this should form the basis of the price which policyholders are charged. This may need adjusting for expenses and other frictional costs. However, it seems unlikely that shareholders and other investors would value an option written on their behalf by a life office significantly differently from the price that they could trade it for in the open market. Further, this should be the case whether the option is hedged or not.

Even if an exact hedge does not exist, it seems reasonable to consider existing prices as a good guide to the value that investors would put on those which are not quoted. In the case considered by the authors, prices exist for products to hedge the interest rate risk and the equity risk, but not the combined risk. However, we can use the option pricing formulae to take the available prices as an input, to suggest a consistent value of a GAO hedge. In this sense, and this is how they are used in practice, the option pricing formulae can be considered a way of 'graduating' available prices. Option pricing techniques should form the basis of insurance pricing in most circumstances, but the authors seem to restrict its use to circumstances where an almost perfect hedge is available.

If we accept that these techniques are appropriate in terms of price and value, we have to ask why offices would do anything but hedge when the opportunity presents itself. The authors make reference to the practical difficulties of hedging, but these seem to overstate the case. They consider only dynamic hedging using government securities, while most GAO hedges would, in practice, be based upon a mix of swaps, swaptions or bespoke over the counter (OTC) products. Markets for the first two products are far larger and far more liquid than the government debt market, while the use of swaptions and bespoke products removes the need for dynamic hedging.

In reality, most offices have now hedged the majority of their with-profits GAO exposures, which is prima facie evidence that it is possible to hedge. Indeed, the market has absorbed this without much turmoil.

It may be that offices do not hedge more fully, because they do not really believe that the value of the guarantees that they have written is as great as the hedging cost; or because to recognise the cost in the balance sheet would be unpalatable; or it may be that hedges are used more often than we are aware of. It will be interesting to see how the new realistic reporting regime of the Financial Services Authority (FSA) changes this approach, if at all.

The authors do us a great service in showing how the appropriate formulae for calculating

such hedging costs can be calculated. While sometimes tedious, the techniques are fairly standard and relatively easily learnt. They are even easier to put into practice, as many of the complex calculations have been addressed in the literature. For with-profits policies the calculations are yet easier, as the correlation between future equity returns and interest rates is relatively unimportant if the bonus declared each year reflects the then current cost of hedging. Such an approach is set out by Pelsser (2002), and leads to the ability to hedge GAOs closely using vanilla swaptions. These techniques deserve more attention from actuaries.

If a hedge is not held, either because it is not available or because the office does not wish to implement it, we must ask two further questions: "How much more capital should be held to absorb potential adverse experience?" and: "What charge should be made for it?" These are the key topics addressed by the authors.

In answering the first question, the authors adopt a 'real world' stochastic modelling approach, based upon the Wilkie model. They consider both a quantile approach, often called the value-at-risk (VaR) approach, and a conditional expected loss, or tail VaR approach, to setting reserves. This second approach sets capital requirements as the expected loss, given that it exceeds the appropriate percentile.

The authors suggest that the tail VaR approach is superior for a number of technical reasons, as is discussed by Artzner *et al.* (1999). I am not sure that I agree with this. The axioms of a coherent risk measure, the main rationale for tail VaR, are designed to ensure that it can be used to decentralise the decision to accept risk. It is not clear why this should be an appropriate basis for decisions at a company-wide level as to the quantum of capital to hold. Rather, it seems that the metric chosen should reflect the concerns of shareholders, policyholders and regulators. Recent history shows that the occurrence of a default would seem to have a disproportionate impact compared to its size, suggesting that their concerns might better be addressed by VaR.

Other views are possible; the main point is that the risk measure chosen should be driven by the stakeholders' views of risk, not by mathematical niceties. On this I would second the authors' plea for greater discussion.

The authors also question at what level capital calculations should be carried out: at policy, product, fund or company level. As capital is fungible, only a fund, or company, perspective makes sense. This requires an allocation of capital for pricing purposes, for which a marginal approach is required. This will almost certainly mean that the sum of the capital requirements will not equal the total capital required. However, any other approach will distort decision making; a different conclusion results if capital is used for risk budgeting, as Artzner *et al.* (1999) discuss.

An issue which the authors do not discuss is the length of time period that any capital projection should cover. Instead, they adopt the approach of the MGWP (Ford *et al.*, 1980), assuming that capital is held until the policies run off the books. This approach has a number of problems. It leads to: results very dependent (as the authors show) on the asset model used; to conceptual problems regarding new business; and to various other issues. In particular, strong mean reversion will reduce the capital required.

Another possibility is to use a shorter time period, followed by a move to a matched position, if circumstances dictate. In this model the time period is that required to notice that a problem has occurred and to hedge it. This approach is common in the banking world, where a ten-day period is often used; the idea being that you can close off anything within ten days. If only the insurance world were so simple! It also appears to underlie the FSA's new realistic reporting regime for with-profits offices.

Where practical, this approach circumvents the normal thorny problems of allowing for new business and of asset model features dominating results. It must allow for correlations between the cost of hedging, following a market movement, and the size of the market movement, which is not something that the FSA's approach yet does. It also does require a reasonably liquid market. It would be suitable for looking at interest rate risk and for market risk, but it would not work for mortality risk, unless we think that the insurance market is liquid.

This approach also leads us firmly to a mark-to-market accounting model for liabilities and capital. The authors prefer marking-to-market where possible, and I strongly agree. Even though the market cost of GAOs was low when they were first written, this would surely have led to an earlier recognition of the problem, hedging and perhaps earlier withdrawal of the guarantee from new business.

The second question which the authors asked is: "How should we charge for capital once we had decided what it should be?" The authors use a fairly arbitrary charge of 1% or 2% p.a. Is there a more scientific way? If we use an option pricing approach to set the basic cost of the guarantee, this has already fully costed the potential loss to the company, even if a hedge is not in place. Option pricing formulae can be derived in an expected utility framework, and such equilibrium models clearly allow for the price of this risk. As such, we only need to consider whether any additional costs arise from simply holding capital within a company; that is that we have a structure question, not a cost of risk question.

The economics literature suggests at least two possible costs, that of double taxation and that arising from the principal/agent problem. These have been a major field of study in financial economics, in parallel with option pricing theory, since the mid-1970s. In my view, the costs of 1% to 2% which the authors have used are probably broadly those required to allow for these items, but the rationale deserves further study. In particular, the Nobel prize winning economist Robert Merton has developed a coherent model of the interaction of such costs, capital requirements and risk management for financial intermediaries, including insurance companies.

On this point, the authors question whether the price of mortality risk can really be dismissed as zero, due to diversifiability arguments. It may be zero cost due to diversification, but the need to hold capital would still give a cost.

Turning to some technical issues, I would have liked to have had more detail on the calibration of the option pricing model, as it is not clear to me that it is truly calibrated to market prices. Similarly, I would have welcomed more discussion around the issue of 'fat tails'. Are they really necessary? How do you model them within a stochastic differential equation framework? For example, the models used for option pricing typically do not have fat tails, and therefore they may not be that important. Models used for capital calculations typically do have fat tails, and therefore are very important. How would the two fit together?

In summary, this is an important paper, as it addresses the key issues which actuaries need to face in relation to pricing and reserving for market risk in the future. I do not agree with all of the authors' conclusions, particularly on the practicality of hedging. However, they suggest a move towards a mark-to-market model that will serve us better than our past approaches, which seemed to have dismissed much of the information available in the markets.

Mr G. D. Clay, F.I.A.: I shall concentrate on the historical perspective in 1985. I returned to the United Kingdom in the middle of 1983 after ten years in South Africa, and have a clear recollection of what an up-to-date with-profits actuary knew at the time. It may help to counter the misperception that the actuarial profession and the life assurance industry were totally unaware of some of the issues mentioned in the paper.

By way of background, South Africa had its market crash in 1969, rather than in 1974, as in the U.K. More relevantly, it had generally higher inflation and lower gilt and semi-gilt rates than the U.K. My office was a major writer of immediate annuity (IA) business in both countries, and we matched, or immunised, our IA assets and liabilities rather closely. Consequently, when I was responsible for introducing a new range of unit-linked endowment assurance products in South Africa in 1974, we did not consider the possibility of including a maturity guarantee. Therefore, the report in 1980 from the MGWP seemed a mathematical statement of the blindingly obvious! Guarantees cost money, and possibly a lot. Fat tails were not publicly discussed then, but the concept was certainly recognised.

In ¶2.2.3 the authors deal with mortality improvement. Although not explicitly allowed for, improvements in mortality were allowed for in the interest rates assumed. Members of the profession were aware before 1985 of the risks associated with both mortality improvement and

the tails in investment performance. I suspect that they would have been surprised at the pace of mortality improvement that we have actually experienced over the last ten or 15 years. However, my life offices saw no particular need to quantify a specific provision until it became an industrywide issue a few years ago, because whatever the basis used, it was clearly small in relation to the explicit contingency reserves held within the valuation liabilities, let alone the implicit margins in the published valuation basis.

Another point relates to the charge for shareholders' capital, factor h in ¶3.2.4. In reality, 1% or 2% was on the low side, even back in the 1980s. From 1987 I was the actuary responsible for all our non-U.K. operations. Our standard approach to the use of capital was to have a hurdle rate for the use of the with-profits estate of the fund earned rate plus 3% p.a. This compared to a hurdle rate for the U.K. life business of the fund earned rate, and for worldwide general insurance business, the expected equity return with an addition in the range of 5%-10% p.a., depending on the riskiness of the business. There was some fairly primitive science underlying these rates. Basically, we looked at the risk-adjusted return on capital that shareholders might expect to achieve; but clearly there was a differentiation between the U.K. and the overseas life businesses, because the estate was a relatively scarce resource overseas and clearly a plentiful one in the U.K. A 3% p.a. target charge for use of capital overseas seemed to us appropriate in the late 1980s. Since then our target returns have moved to being based on the expected equity return rather than on the fund earned rate, while the risk margins required have continued to evolve.

I trust that this scene setting will be useful in the discussion.

Professor P. Artzner (a visitor)**:** This paper is most welcome at the time of Solvency II and of changes in international accounting standards. I will make two observations on two topics from the title, reserving and pricing.

In Section 3.4 the title is *Reserving from Year to Year*. In ¶3.4.2 it is stated that, in 1980 the MGWP "felt ... that it was desirable ... that there should ... be no requirement for additional reserves during the currency of the policy." So, reserves had to do with risk and its measurement, and maybe with its management. From a background of multi-period risk measurement, I am looking for some sort of valuation, or appraisal, of a policy as opposed to mainstream financial mathematics.

There may be room for risk-adjusted appraisal, and I will identify two reserves. Higher risk-adjusted appraisal should be like higher reserves. Then I will also identify what supervisory actions should be taken once reserves have been properly defined.

Firstly, appraisal of policies requires a valuation with a retrospective approach for the flows which have already been seen, and some view of the future, like technical provisions. An appraisal would take into account the premium, maybe a single premium paid at time zero, investment return and pensions at time t. Assuming no intermediate claim, and that the guarantee comes at the end, and the valuations $X(0), X(1), \ldots X(N)$ of the policy on the part of the insurance company, we have risk-adjusted appraisals $\tilde{X}(0), \tilde{X}(1), \ldots \tilde{X}(t), \ldots \tilde{X}(N-1), (N-1)$ being the last date before final termination of the policy. The $\tilde{X}(t)$ have to be conservative estimates of the value of the policy to the company.

A supervisor, internal or external, may follow the rule that, before letting the business go, he requires that the risk-adjusted appraisal at $\tilde{X}(0)$ be positive. If we go back to the reported opinion of the MGWP, the initial reserves should not be changed over the year. I see that as requiring a bit more. For this we define the sigma of X, the first date at which the appraisal becomes negative, where there is bad news. Then I understand that the suggestion made, which was not the one retained in the paper, is to require that at any date before the date sigma of X, the risk-adjusted value be positive. It is some formalisation of the suggestion in ¶3.4.2. There it seems very natural, but there are some technical difficulties.

For example, with the prescribed intervention policy it is possible that a process of appraisal values may be accepted by the supervisor, either internal or external, while, at the same time, a richer appraisal process where everything is ten units, would be rejected, which is strange. For example, consider the 'Tree' in Figure A. The values above the nodes, in bold type, refer to the X

https://doi.org/10.1017/S1357321700004578 Published online by Cambridge University Press

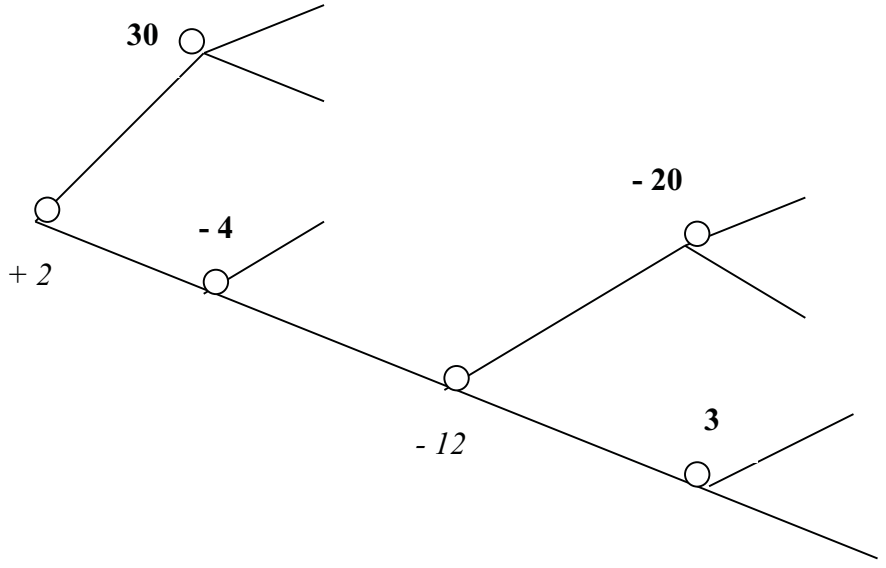

Figure A. Tree example

process, the appraisal value. The italicised values below the nodes refer to the risk-adjusted values X. At time zero you find that everything works well: the risk-adjusted value is positive, equal, say, to '*2*'. Why? Because there is a large 30 or a small −4 at the next date. We are in trouble here, but there is only a low probability of being in trouble. Since you cannot hope to have no 'trouble at all', you would allow it to go at time zero.

If time one is against you, you are in trouble. In the first formulation everything stops there, and you do not see the terrible −20, which comes a few periods later. So, the negative italic numbers of risk-adjusted values are irrelevant, because at time zero you know that at time 1 there is a possible insolvency, and you reject the business. However, you are ready to accept the situation where there is trouble at some future time and the risk-adjusted value is positive at the time before. The surprising fact is that, if the contract gives ten more units to the company for any time, you would reject it.

The big trouble at date three, −20, is a medium trouble now, only −10. However, the warning bell one period before the risk-adjusted value, which serves as an alarm to tell you whether to stop the business or not, being at −12, plus ten more units, is negative. So, with this intervention policy at the initial date, the supervisor, either internal or external, would say: "Sorry, we cannot let this business go. There will be one point where are you are insolvent and you would have been warned of this before. We do not let you go."

The mathematical explanation of this is that there is a complicated connection between the first time of insolvency for a policy X and the first time of insolvency for a policy X plus ten.

In ¶¶2.3.10 and 2.3.11 the paper quotes a paper of Professor Wilkie (Wilkie, 1978), where a suggestion is made to price a policy out of expected claims plus compensation due to shareholders bringing some risk capital. That is very interesting as a connection of risk measurement, as the capital required gives us a price for policies. It puts some light on some current debates, or battles, about the relationships between risk measurement and insurance premium principles.

More importantly, at a time of Solvency II, this consideration about pricing brings some light on the discussion between solvency being provided by larger solvency margins or by more conservative provisions.

Mr D. B. Duval, F.I.A.: I want to focus on the lag time between introducing a new theory and assimilating it into practice, and how long it is and/or it should be. There is an implicit assumption in the paper that, in 1985, it would have been a good idea if all life office actuaries had picked up the theories in the papers as then published, and used them immediately. Clearly, that would not, in fact, have been a good idea. Most of them would not have understood what they were doing, because it takes time to absorb and to fully understand a new theory. Secondly, there are issues in the theory which are quite definite traps. The obvious one is the point made in the paper about quantile risk reserving being incoherent, and therefore, if you do it wrong and add up your reserves, you can actually end up with a reserve that is too small. I recall that this was pointed out at the time, that if you use ruin probabilities and amounts of risk of ruin, you can get rather implausible answers.

To assume that a new theory should immediately go into all professional practice is inappropriate. Clearly, it should be made available, as soon as possible, so that people who want, and are able, to use it for commercial advantage can do so. The question for the profession is: "At what point should we insist that it is used in our professional work?" Actuaries are primarily giving information to other people, on which other people make decisions. Actuaries, themselves, make very few decisions, but we give a great deal of advice to boards, or information to regulators, which may lead them to make particular decisions. In fact, we appear not to have a structured approach in deciding at what point something should be coming into practice. The time period should be long. With other professions, like medicine, for example, it is quite a long time from when you see a new medical technique described in professional journals to it becoming the standard treatment. We do not know at what point a comparable situation arises for actuaries. Is it, for example, essential that stochastic modelling, or option pricing, should be applied to all significant financial risks, and actuaries reporting on such an institution should include such things in their reports, which is not the case at the moment? There are a number of sizeable financial institutions, insurance companies and pension funds, where the boards have not received such advice, and the profession is currently silent as to whether we think that it is appropriate or not.

Actuaries have given advice on the risks being taken through scenario testing, but not in the particular framework of stochastic modelling or option pricing. Do we, as a profession, think that this is acceptable, or not? The answer at the moment, presumably, is, since we know what we are doing and no-one is objecting, we believe that it is acceptable. I find that a very uncomfortable position. The only way in which we seem to change standards absolutely is through the examination system, which is an incredibly slow process. Consequently, other people are setting our standards. The majority of the setting of risk standards and the setting of fair value standards is now done by the regulators and the accountancy profession. We simply calculate the numbers in accordance with the principles that they have laid down, or sometimes we work out the interpretation of the principles. Whether they are the correct principles or not, we may well seek to influence, and we do very actively, but we do not determine them. We have effectively abdicated that part of our professional responsibility. That is a great pity, and, looking forward, we need to find a way to reverse this trend. For example, VaR is a very widely used measure in the banking sector. Despite the opener's comments, its flaws are significant, and conditional tail expectation is a far superior measure. If the majority of the profession agree with me, it would be better if that opinion were published.

In summary, there are two things which we need to do, which are both essentially jobs for the Professional Affairs Board. The first one is to look at what practising actuaries are not doing now, but which, according to theory, they should be. The second is the question of basic minimum standards. What minimum standards do we expect in actuarial work in the various practice areas, and are they consistent?

The President (Mr J. Goford, F.I.A.): When the Government Actuary's Department wanted to introduce the resilience test, it brought a paper to the Institute. The points in the discussion for and against the test were counted, and afterwards it was deemed to have been passed into legislation. I think that there were a few objections at the time.

Mr W. M. Abbott, F.I.A.: Mr Duval hinted at a 'new' work programme for the Professional Affairs Board. There has been talk of an Actuarial Standards Board, and much of the technical content would be far better incorporated in the guidance of a semi-independent, or even independent, Actuarial Standards Board than within guidance issued by the profession, with the aim of making changes more rapidly.

Turning to the paper, the last part of ¶12.8 needs to be clarified. This says: "the option price is only a guide to the 'fair value' if the required hedging strategy can practicably be followed and then is actually followed." The last five words are unnecessary, because what the company actually does does not matter in accounting terms, so long as it makes an effort to hedge the liability portfolio. If the required hedging strategy could be done, but is not done, then that is a case for the regulator to demand an additional capital reserve. If the required hedging strategy is impractical, then the liability should reflect a risk margin to cover the mismatch risk from the nearest practical investment strategy.

Some of the language in ¶12.11, which talks about a contingency reserve, is probably not right, since, in the accounting principles which are now being developed, contingency reserves would not be part of a liability.

Professor L. C. G. Rogers, Hon. F.I.A.: It is clear from the paper that there are competing ways in which you could try to value contingent claims, and value policies, which would have some sort of random payout at the end of the policy. It is clear that the basic methodology, the idea of taking the expectation of the payout, cannot be quite right. For example, there is a share which is trading today for £100, with an annualised rate of return of 20% and a volatility of 20%, and the riskless rate is 5%. So, in a year from now you are going to have something which is log-normally distributed with those parameters. What is the fair price to pay for that share today? It is not the expected value of what it is in a year's time, but the fair price to pay for it today, which is £100. When there exists a hedge, the fair price for the contingent claim has to be the price of that hedge. That may be the case in an almost complete market, such as you would find in investment banking, but it is very far from being the case in insurance, where the market is incomplete. You are facing the threat of viruses like AIDS, for example, which may suddenly quadruple the cost of life insurance policies. So, how would you price something in an incomplete market? The most extreme example, where there is absolutely no hedge available, is a toss of a coin.

If I toss a coin and it comes down 'heads', you can have £20,000, and if it comes down 'tails' you can have zero. Contrast that with a gamble where you will certainly get £10,000, and we would all prefer the latter. This illustrates the simple economic concept of risk aversion, that people would prefer to have a certain sum of money than a random sum of money with the same expected value. So, somehow you have to pay for risk and be compensated for taking on risk.

In my example, where there is absolutely no hedging instrument, how would you decide at what price you would be prepared to take it? You were not prepared to gamble 50-50, zero to £20,000 for £10,000, but, if it was laid against a certain £9,000, would you do it? If it was laid against a certain £8,000, would you do it? If I bring the price down, eventually there is a price at which you will take that gamble. It will be a different price for everybody, depending on their different appetites for risk. The situation that we are considering is one where we have to be able to answer questions concerning what price we are prepared to pay for taking on some risky gamble. Not every risk can be hedged, but we still have to be able to answer that question.

If you were to present this dilemma to a group of economists, they would say that this is well understood. We look at the expected utility of your reward, so each individual agent has his utility. If he did not take the gamble, he is going to receive a certain pile of money at time one. If

he did take the gamble, he is going to receive that pile of money plus some random amount. The price that he is prepared to pay for that random amount is the price that would exactly equalise his expected utility. If he were to receive that random amount and pay £5 and get exactly the same expected utility as if he did not enter into the gamble, then £5 is the price that he is prepared to pay for it.

The conceptual framework for handling risks is well understood, but in practice is a difficult issue, because you are dealing with, say, the performance of a large company, where you are not talking about just one option or one GAO, but about the entire portfolio of the company. You are also thinking about time frames and cash flow, as time progresses, so you have to be more intricate and computationally intensive. If you can specify the distributions of the cash flows that you are proposing to take on, then you can in principle work out: (a) what price you will be prepared to pay for them; and (b) what dynamic hedging position you would take as a function of the prevailing conditions at all times between now and the expiry of the option.

The investment banking industry has powerful tools around that can handle these kinds of computations, albeit not necessarily very well.

Mr D. I. W. Reynolds, F.I.A.: I sit on the Professional Affairs Board, and am pleased to inform Mr Duval that we are considering some of the issues which he raised, in particular the difficulties that may be encountered by the profession when changing its guidance and standards. For example, if you change your professional standards you could find a lawyer claiming that your past practice was wrong, and therefore you are liable in one way or another.

One aspect of the paper gave me some concern, which is the issue of the market price and the question of whether there is a single market price. VaR is based on market price, but if a large investor, such as the Prudential, owns 5% of a stock, is the market price of that equal to the value for which you could sell it? Clearly, the answer is 'no'.

I have a problem with too much theory being based on market price, which links up with Professor Rogers' comments concerning a different market price for different people, versus my query that there is actually a different market price, depending on how big the holding is that you have. There is a lot of good value in financial economics, but I worry that, on the basis of using market price and volatility as the measure of risk, it is actually building on uncertain foundations.

Mr B. Bergman, F.I.A.: Until a few years ago I worked for a company which was a market leader in unit-linked GAOs. When the GAO problem came to prominence, particularly in 1996, and probably up until 1999, when falling interest rates really caused problems, we saw large increases in reserves as interest rates fell. We wrestled with a number of approaches to containing these spiralling reserve increases. Initially we attempted dynamic hedging. During the 1998 Russian crisis I lost faith in dynamic hedging, as I could see how one could actually shift the market as one traded one's assets. We then moved on to swaptions, but these escalated from the time of first inquiries in early 1997 to the time when we actually considered purchasing, by which time they had escalated beyond belief, due to volatility changes, which were totally in the control of the investment banks.

However, a point not mentioned in the paper is that, irrespective of whether you use the approach of the MGWP (Ford *et al.*, 1980) or a financial economics approach, there are some issues that are not mathematical, but which are equally, if not more, important. These relate to policyholders' behaviour, and to the relevant rules and regulations in place.

I have a list of four issues relating to GAOs which could swamp any different mathematical models which you might consider. Firstly, tax-free cash and the assumption that one makes when deciding upon the question of liabilities to hedge could make a difference of 25% to the reserves held. Policyholders do not always act efficiently. It is difficult to buy a car with a pension, but you can buy a car outright or go on holiday with tax-free cash. Second, there are trivial annuities and the legal permitted threshold at which a lump sum can be commuted for cash. Every time when a threshold was increased millions of pounds of liabilities 'disappeared'.

Thirdly, some policies only gave the guarantee if the policyholder took a certain type of annuity. Frequently, these were single life annuities payable annually in arrears. One's company might have taken a more generous view to that, but where does one pitch one's hedge when one is uncertain about the actual number of policyholders who are going to opt for a single life annually in arrears annuity? Fourthly, there are concerns about specific retirement ages or dates. The paper refers to age 65, but we had policies which only gave guarantees on certain dates. Again, a more generous strategy could have been adopted, but what difference does it make what model you use when the assumption as to when the policyholder is actually going to retire is so critical?

These four issues swamp the niceties of the mathematical modelling considerations.

The President (Mr J. Goford, F.I.A.): I think that we are currently somewhere between reserving purely for what is in the contract and reserving for expectations, reasonable and unreasonable. We are definitely moving in the direction of reserving for expectations. Whether the additional reserves are a provision or are in the capital requirements is debatable. I think that, whenever we think about reserving for what is absolutely in the contracts, then there are certainly some people outside who might disagree. So, we have to be careful about whether we should be reserving also for expectations.

Mr A. J. Wells, F.I.A.: My first comment concerns the practical difficulties of hedging. Mr Reynolds raised the example of the Prudential, perhaps owning 5% of a particular stock. He queried whether the value that it could place on that stock was five times the value that a holder of 1% of that stock might place on it. A different approach is that of a very small life office, which might have a complicated portfolio of GAOs, worth, say, half a million pounds. Could it get an over-the-counter option to back those at an efficient price versus a very large life office which might have a similar equally complex portfolio of GAOs, but 50 or 100 times the size? I am sure that an investment bank might be happy to construct such a high value portfolio, but could the market bear it, and would the cost to that large office of hedging its GAOs be 100 times the cost of the small office?

On a completely different point, there has been discussion this evening on the lag between theory and implementation in practice. Academics might develop new theoretical ways of doing something, which might be discussed by the profession, with some actuaries starting to use it, and many years down the line it might become generally accepted practice. However, legislation and guidance talk of offices using generally accepted actuarial methods, principles and techniques. An individual Appointed Actuary has to consider when a modification to a technique, or a new technique, has become generally accepted. One can imagine the poor bedevilled Appointed Actuary thinking: "I have to convince my colleagues, the Board of Directors, the profession and the FSA. It may be too big a job for this year.''

If one Appointed Actuary thinks like that, the office might be a year behind the times. If all Appointed Actuaries think like that this year, next year and every year, there could be a further significant delay before a new technique becomes generally accepted.

Mr P. J. Tuley, F.I.A. (closing the discussion)**:** While the practicalities of modelling are not lightly to be dismissed, one should not forget that the option we are talking about was heavily dependent on longevity, a danger that itself has only recently taken centre stage. There are many who are smarting from that. There are numerous blocks of immediate annuity business that, with hindsight, are unprofitable, just as these GAOs are.

Perhaps, when looking back over history, what counts are the risks that were weighing on the minds of those overseeing the firms at the time. One lesson of the paper is not to be too focused in one direction. That is where discussion with those, outside the day-to-day mill and grind of life, can be very valuable. True risk assessment, trying to think of all the risks that can bite, is actually very scary for most of us, conditioned by our histories.

When one talks of models, one often refers back to the MGWP (Ford *et al.*, 1980). Again, it

is of interest that the current models being brought in for with-profits business for maturity guarantees are only now moving from development into production. They have cost various firms a great deal of money to get into practice. Bringing these techniques into actual production is no easy matter.

The paper is looking at GAOs in a unit-linked context. It is, perhaps, that that leads the authors to conclude, in ¶10.6.11, that hedging is unlikely to be useful. In practice, that arises from the practical problems of unit-linked GAOs; namely, that the reserve for GAOs is partially forced to be invested in units itself under the close matching provisions of our legislation. That is unfortunate, and makes the quanto nature of the option (a fluctuating nominal onto an option of varying worth) far worse. By contrast, numerous with-profits GAOs have been hedged quite successfully.

The paper then considers the mark-to-market of capital. The thrust of regulatory development is to mark-to-market for the basic reserve and to move that way for the capital. That thrust is going to be with us under a number of guises, from a number of different regulatory or accounting sources. The prudent statutory reserving basis might effectively use a different system, a mark-to-market system on triggers, where you reassess the reserve only after particularly large moves in what the market reserve might need to be. It is the recent history of both mortality and interest rates that make the description in the paper show that method to be so lagging and out of date. However, it is not a change which we should ignore, because, while the paper is firmly stochastic, it is not necessarily close to market prices. You may get similar lagging effects by using historic assessments of volatility, so that any long trends in volatility are only slowly recognised through the reserving system.

That is really the problem. One is in a model world, not the real world. If you subsume to market consistency, or inconsistency, you may have an open market, or a very thin market. However, these numbers drive real management actions. They actually impact claim values. So, the robustness of the model results is very important and is a big issue for us. For that reason, I strongly support current market value as the driver of these numbers, because that is a firm base, which can drive management actions to hedge out some of these risks.

To add to the list of things that we could explore, I would support conditional tail expectations, which has been touted in North America and in international circles, because it sits very well with a desire to close out risks if the size of the failure that you are actually worrying about is uncertain. Indeed, if we look at that we might even be ahead of our banking colleagues!

Professor H. R. Waters, F.I.A., F.F.A. (replying)**:** One issue discussed, starting with the opener's remarks, was the practicalities of doing the hedging. The opener implied that the authors had been a little reticent, and that it was not as difficult as we set out in the paper to do the hedging, and many offices had done this. That is gratifying to hear.

However, Mr Bergman said that he had tried to implement dynamic hedging from 1996 onwards, and found that it shifted the market, and there were various practical difficulties which prevented him from doing this. This was in the context of a unit-linked office. Another speaker said that hedging was more likely to be successful in relation to a with-profits office. I am not sure why it would be more difficult in relation to a unit-linked office than to a with-profits office. This takes us a little off the topic of the paper, because we were just trying to set out a methodology.

Another point that the opener made was that, since the liability can be hedged, why would an office do anything but hedge it? Financial economics does not provide all the answers. What happens in our paper is that we take a financial economics approach in Appendix C and in Section 10, and say: "Let us imagine that the market consists of three tradeable assets, whose prices behave in very particular ways. They follow some very particular models, for example geometric Brownian motion, and we know that they are going to follow precisely those models with the parameter values that we assigned to them." Given that we accept all of this, we can then set up a perfect hedge, provided that we can trade continuously and costlessly. At this stage, of course, we are in a model world, and we recognise that we do not have all the answers.

What we do get from the financial economics approach is an idea of a dynamic investment strategy that may well help us to manage the risk, not reduce the risk to nothing, but to manage the risk.

A good analogy here would be Redington's theory of immunisation, that is, if you have a level yield curve and certain fixed liabilities, and the yield curve makes a level shift of a small amount, then, by immunising your setting up of the proper investment strategy, you can make a certain profit, as you will be immunised against any losses. That is working in an ideal world, where you know the liabilities, where you have level yield curves, and any shift in the market is to another level yield curve, and that shift can only be a small one. The same is true when we are talking about financial economics. It is giving us an insight as to the dynamic investment strategy that we could be following to reduce our risk.

Mr Clay's historical perspective was interesting. Working in South Africa in the late 1960s, with the market crash there, he had an advantage over people working in Britain in the 1970s. I agree with him that actuaries were, and are, aware of risks inherent in interest rate movements and in mortality. This is bread and butter to actuaries, which, perhaps, makes it all the more surprising that GAO risks, which are just a combination of interest rate and mortality risk, were not discussed more within the profession. Professor Artzner, Professor Wilkie and I have discussed ideas on multiperiod risk measures many times, and that discussion will go on, because there is no definite answer to the problems that Professor Artzner was outlining. One of the general points that I picked up from what Professor Artzner was saying was that the MGWP (Ford *et al.*, 1980) was trying to set up a reserve so that there would be no requirement for additional reserves to be put into the portfolio at some future time.

That ties in very well with my actuarial education, which is that, if you have a future liability, you should set up the reserve for it now. That is not the way in which things have developed. We now have marking-to-market, where you are going to release reserves, and strengthen reserves as the market moves. There has been a lot of agreement here that marking-to-market is a good way to proceed, and there was general agreement in the Faculty discussion about that. The ideas of the MGWP are being left behind.

I was interested in the point made by Mr Duval about how a profession moves forward in its thinking, which is an interesting topic for discussion. We might even benefit from talking to other professions at maybe even the sharper end of practice, for example the medics could have something to say to us.

Mr Duval implied that there are several institutions, insurance companies and pension funds, where stochastic modelling is not carried out. In Section 12 we quote from Guidance Note No 8, where it says that if you have liabilities that are stochastic in their nature, then you should be using stochastic methods.

Professor Rogers talked about a fair price being the price of setting up a hedge for any contingent claim, if such a hedge can be set up. In the case of insurance, all too often, where we have an incomplete market, the hedge does not exist. The economists, as he pointed out to us, have ways of dealing with these things, using expected utility, and so on. We have a certain amount to learn there, and we look forward to the further involvement of Professor Rogers with the profession.

Mr Wells talked about the practical difficulties of hedging. He also talked about the difficulty of a small life office, with some complicated liability, buying an over-the-counter option to reduce its risk, and what price would it get compared to a large life office, and so on. One of the points that we make in the paper is that what we try to do is to look at how a life office could actually hedge this risk itself internally. An alternative is for a life office, or for another institution, to go to an investment bank and to purchase some over-the-counter security that will reduce its risk. I am sure that that is what many institutions are doing. That then raises the question of how that investment bank, or that counter party, will actually hedge the risk. It may well be that it has an advantage here. It will be working in a different regulatory framework, so that it will not have close matching requirements, and so on, so that it can do things that the primary life office could not. However, we need to be careful to be sure that these

people are hedging the risk themselves in an appropriate way. It is completely analogous to a primary office buying reinsurance. The reinsurance is only as good as the reinsurer.

The closer said that we, the authors, claim that hedging was unlikely to be useful. That is overstating things, as we did say that there are some practical difficulties. Looking at financial economics and the hedging strategy gives you some ideas as to where you should be investing, and how you should be investing your money. Therefore, I am not sure that I completely agree with him. He also made a comment about when you look at the whole risk assessment issue and the full picture is very scary, and I believe that.

That leads me to make the comment that our paper leaves many questions unanswered, and the biggest question that it leaves unanswered is: "How do you actively manage the risk associated with a portfolio of these types of policies?'' The opener at the Faculty discussion, Mr Speed, made the comment that we have, with GAOs, a liability which becomes greater on the office if equity prices go up, because of the quanto effect, and it becomes greater if the interest rates go down. It may be that there are parts of the operation of the office where you can balance things here against these risks. That gives you an opportunity to look at the risk management on a more global scale for the office, which is something beyond the scope of our paper.

The President (Mr J. Goford, F.I.A.): This is a critical and important paper, because it is part of the process of listening to academics, as well as to lawyers and to the press, to ensure that we move with the mores of society and take into account the expectations of our end customers.

I note Mr Duval's road map, and we will certainly think about the issue that he raised, concerning when it is that something becomes generally accepted by, or within, the profession.

I would not like it to be thought that with these papers, and the changing mores, we are always looking for a generally accepted rule, which is then applied. We will still need to exercise actuarial judgement in applying them to particular issues and problems.

I express my thanks, and the thanks of us all, to the authors, to the opener and the closer, and to all those who have participated in the discussion.

WRITTEN CONTRIBUTIONS

Mr D. M. Pike, F.I.A.: Following my comment at the Faculty discussion of the paper, Professor Wilkie sent me a Word version of the GAO paper, from which I have drawn a few graphs. What I have learnt from the exercise is that, for the 'serious' reader, the tables show much more information in a more concise form, even if graphs can convey the salient points more quickly to the reader who has less time to spend. Here are some examples. Figure D.1 illustrates points (f) and (h) of ¶4.2.2, Figure D.2 illustrates point (e), the yield curves of C99,1, and Figure D.3 illustrates the means, also from point (e). The means show both rising 'yield curves' and (at six years with the lowest interest rates) falling 'yield curves', but the feature diminishes at higher quartiles, so that Q99.9 always has a rising 'yield curve'. The 'increase first and reduce later' feature is mentioned in 4.2.2(e), and is only shown by the policyholder charges, as illustrated in Figure D.2, and again the feature diminishes at higher quartiles, so that C99.9,2 is always rising.

Using ¶5.2, I produced Figure D.4, which shows some of the measures over time, to illustrate how the increases at term ten are proportionately greater than those at term 40, and how the increases for the mean are proportionately greater than for Q99. As stated in the paper, "the same remarks are true for other entry years'', but the feature does appear to diminish at years with lower interest rates, especially at term ten. Intuitively, I would have expected mortality to have a proportionately greater effect at longer terms and at lower interest rates.

Figure D.5 is similar to Figure D.4, and illustrates how this time the increases are greater at term 40 than at term ten, but, as before, are proportionately greater at the mean than at Q99. Again the increases appear to be lower at lower interest rates, more so at shorter terms.

In ¶6.2, which is illustrated by Figure D.6, it is stated that "the new model sometimes

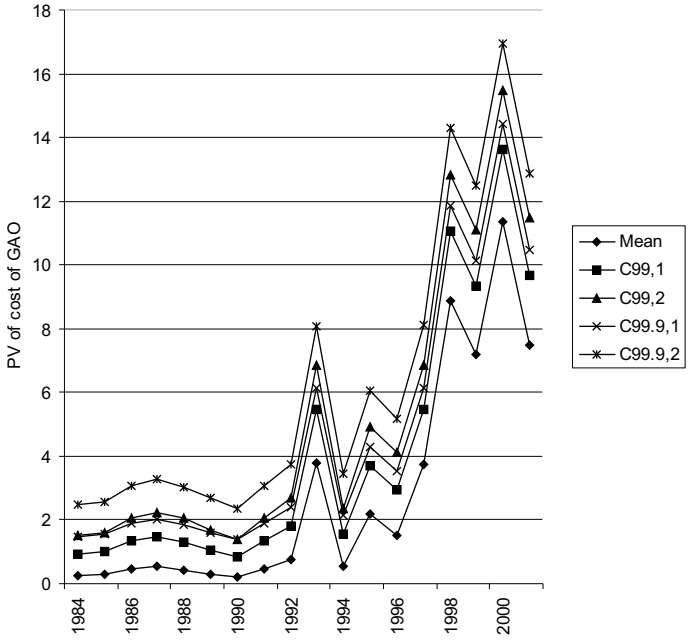

Figure D.1. Illustration of ¶4.2.2(f) and (h) and the results of Tables 4.2 for term ten, of Wilkie *et al.* (2003)

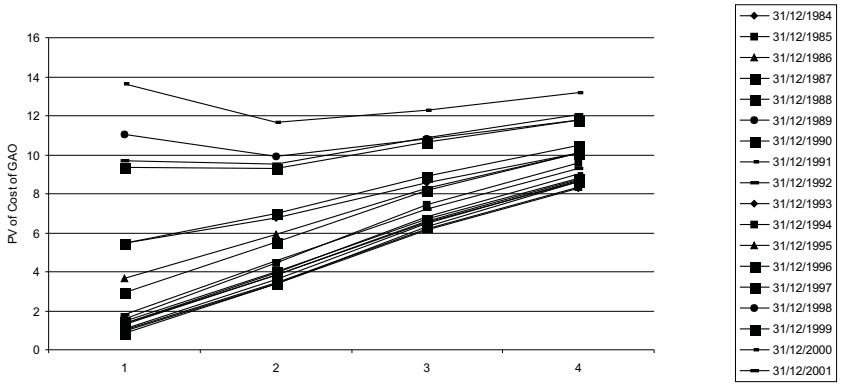

Figure D.2. Illustration of ¶4.2.2(e) of Wilkie *et al.* (2003) for the 'yield curves' of C99,1

https://doi.org/10.1017/S1357321700004578 Published online by Cambridge University Press

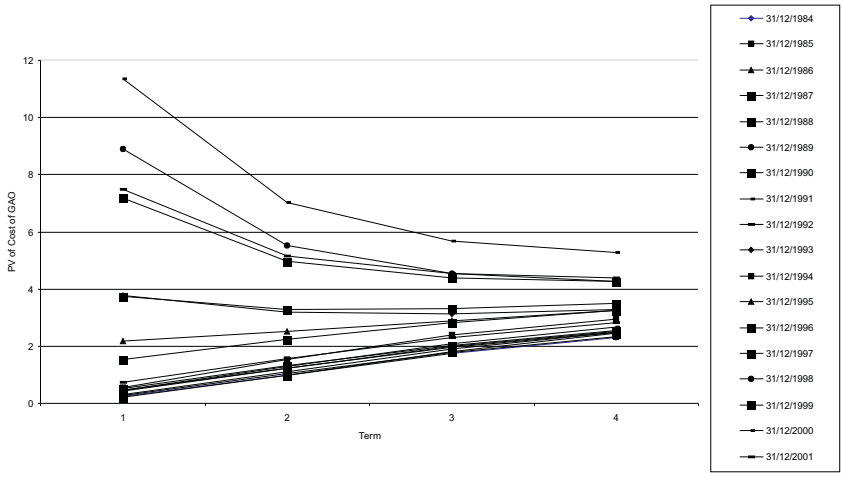

Figure D.3. Illustration of ¶4.2.2(e) of Wilkie *et al.* (2003) for the means

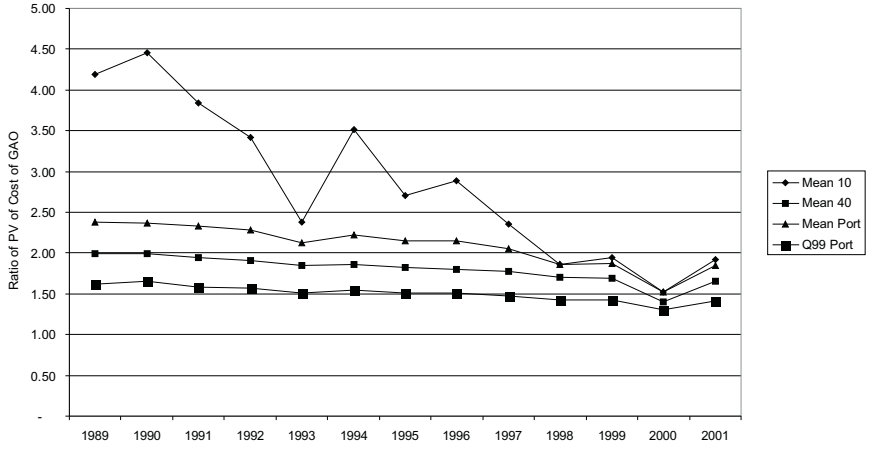

Figure D.4. Illustration of ¶5.2 of Wilkie *et al.* (2003), showing the increases for PMA80 over PMA68

https://doi.org/10.1017/S1357321700004578 Published online by Cambridge University Press

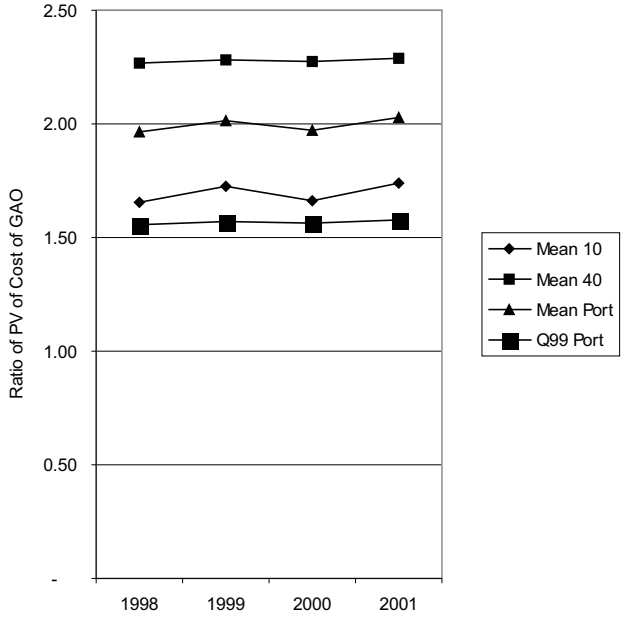

Figure D.5. Illustration of ¶5.3 of Wilkie *et al.* (2003), showing the increases for PMA92 over PMA80

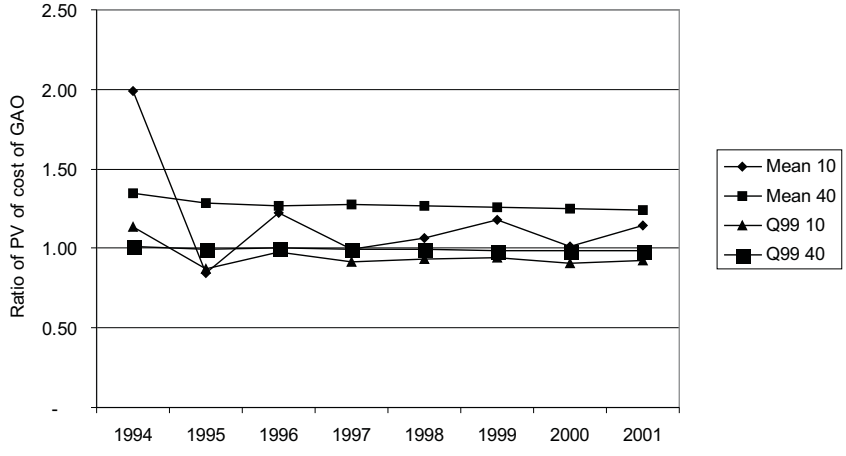

Figure D.6. Illustration of ¶6.2 of Wilkie *et al.* (2003), showing increases using Wilkie 1995 model compared with the 1984 model

https://doi.org/10.1017/S1357321700004578 Published online by Cambridge University Press

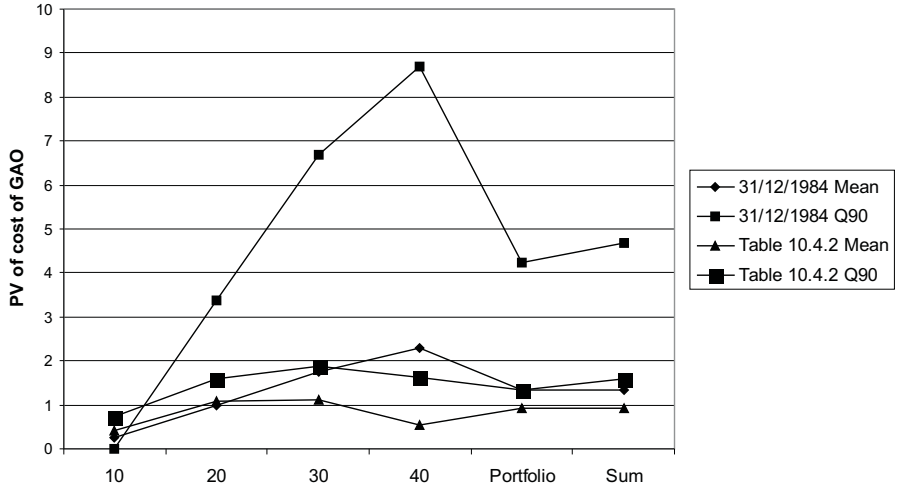

Figure D.7. Illustration of ¶10.4.2 of Wilkie *et al.* (2003), hedging as of 31 December 1984

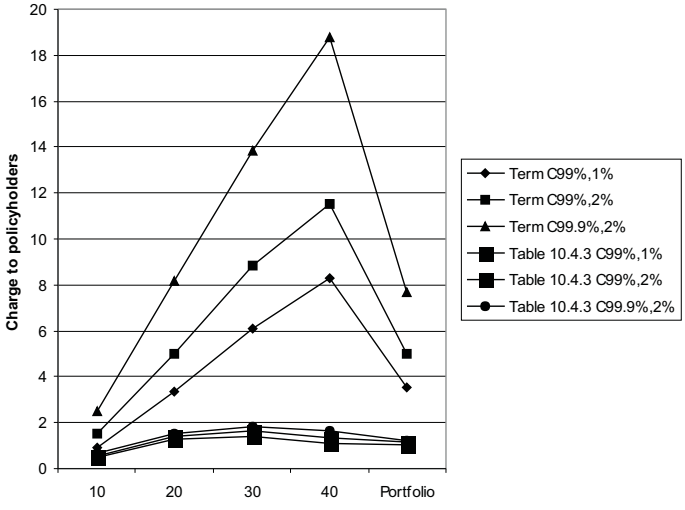

Figure D.8. Illustration of ¶10.4.3 of Wilkie *et al.* (2003), hedging as at 31 December 1984

produces higher values, sometimes lower.'' However, means are generally higher (except for 1995), more so at higher terms (except 1994). Higher quartiles are slightly lower at short terms (except for 1994), and very close at long terms. The 1995 model tends to be higher at lower interest rates, at least at short terms.

Figures D.7 and D.8 attempt to illustrate the points made in ¶¶10.4.2 and 10.4.3.

https://doi.org/10.1017/S1357321700004578 Published online by Cambridge University Press

The authors subsequently wrote: The opener claims that "it seems unlikely that shareholders and other investors would value an option written on their behalf by a life office significantly differently from the price that they could trade it for in the open market. Further, this should be the case whether the option is hedged or not."

This raises two important issues. First, GAOs, like most other life office contracts, are not, and cannot be, traded in the open market. Only authorised life offices are permitted to write assurances and annuities on human life. No life office trades either assurances or annuities in the sense of being equally willing to sell them and buy them, and quoting bid and offer prices for them. Indeed, the adverse selection aspects of this, in the sense that the state of health of the life assured has a substantial influence on the price, would make such a market impractical. Arguments from financial economics about what happens on traded markets therefore have restricted applicability to life assurance.

On the other hand, our hypothetical 'deferred life annuity', being a loan repaid by a fixed schedule of decreasing payments, comprising both income and capital, could exist in a traded market, but, so far as we know, few such contracts exist, and none are deferred in the way required to match GAOs.

Secondly, our approach to valuing an insurance liability is to see what the available assets are that can best match that liability, and then see what their market values are. If they were to match almost perfectly, as a dedicated bond portfolio could match a portfolio of immediate annuities (subject to mortality uncertainty), then the value of the annuity portfolio is the same as the value of the bond portfolio. If the office chooses not to invest in the bonds, it then needs to set up, for its own satisfaction, even if not to satisfy any supervisors, a suitable mismatching reserve. However, mortality uncertainty should require a larger bond portfolio, just in case mortality improves more than expected, and contingency reserves are also needed to cover this possibility.

It is the same with a liability that contains an embedded option. If it can be matched by a suitable hedging strategy, then the value of that hedge can be taken as the value of the option. If it cannot be hedged perfectly, and no hedge in any market can ever be perfect, then contingency reserves to cover the hedging error are needed. This leads back to our three values: best estimate, prudent reserve, and fair value. To talk of one value for any insurance liability is almost always misleading. If the office does not intend to hedge, then larger contingency reserves may well be required, and the corresponding best estimate and fair value may be altered. However, if another hedging strategy would produce a lower fair value, then probably that lower value should be taken as representing the fair value required for accounting purposes, because another life office could presumably adopt this better hedging strategy and would be willing to take over the portfolio for this lower amount.

Thus, best estimate and fair value depend on what the best matching investment policy would be, whereas the required contingency reserve depends on what the investment policy of the office actually is. Further, fair value and contingency reserves depend on what the security level required by the market or by the supervisor is, and on what price the market puts on the risk, i.e. the discount rate h, j or k, as used in the paper. We consider this to be a more realistic approach than the theoretical (but impractical) idea of there being a unique market value for everything.

All three values, in addition, may vary according to which models are used for the real world simulation and the hedging process. Thus 'fair value' is not capable of objective determination. The same problem occurs with any method of estimating the 'market value' of any asset or liability that is not traded. The value of a property depends on some surveyor's estimate. The value of an annuity portfolio depends, not only on interest rates, which can be determined reasonably objectively, but also on some actuary's estimate of the mortality rates that will apply to that portfolio over its future lifetime.

The opener then suggests that we seem to restrict the use of option pricing techniques to circumstances where a perfect hedge is available. This not correct. Option pricing may provide a hedging strategy, an investment policy, which needs to be tried out within a (necessarily

simulated) real world model. An office may wish to test out any investment strategy, whether a complicated one or a naïve one, within its simulated 'real world' model. A better option hedging model should give better results (in the sense of lower contingency reserves, and lower fair value) in the simulated real world, and therefore probably better results in the real real world. We support using whatever good models offices can find.

Several speakers, including the opener, suggest using swaps, swaptions and bespoke products rather than relying, as we do, on government securities. Such products are, in our view, quite permissible ways of covering the liability, but they do involve some counterparty risk. We discuss this in Section 10.7 on *Reinsurance*. The counterparty requires to hedge, ultimately using basic investments, not other derivatives. If the only long enough bonds available are commercial bonds, then the security is less than with government securities, and the risk of default a long way ahead cannot be treated as negligible. It is interesting to study Macaulay (1938) who uses railroad bonds as having the highest credit rating and being free of many other complications, such as being able to be called, which United States federal bonds had at the time. It would be interesting to investigate how many of the over 70 railroad bonds outstanding in 1938 that Macaulay lists, some still dated well into the future in 2004, were repaid in full or are still extant. (For example, a certificate for $10,000 of West Shore Railroad 4% First Mortgage Guaranteed Bond, issued in 1937 and repayable in 2361, can currently be purchased on the internet for $22.00.)

The opener notes that Pellser's method allows GAOs to be hedged closely using vanilla swaptions. We discuss Pellser's paper in ¶¶10.5.5 to 10.5.7. We agree that Pellser's method suits reversionary bonus with-profits policies, but it leaves any variable terminal bonus uncovered, and there remains counterparty risk.

The opener is not sure whether VaR (quantile reserves) or tail VaR (CTEs) is preferable. We do not hold either view strongly, though we rather favour CTEs; but the 'incoherence' aspects of VaR should be remembered, in particular that a quantile reserve in some circumstances can be less than the mean, and that the VaR of a portfolio can be greater than the sum of the VaRs of the individual components. Provided that these unusual circumstances do not apply, either works, but provides different security levels if the same value of α is used.

The opener refers to the question of "at what level capital calculations should be carried out, at policy, product, fund or company level." We entirely agree that the company level is of overall importance, but, for practical calculation it may be easier to calculate at a lower level, even at policy level, with a lower security probability, and then aggregate, after trial investigations as to the appropriate probability to choose, as we discuss in Section 3.3.

We have used a run-off time horizon, whereas VaR is often based on a shorter period, which allows for the time taken to achieve a better hedged position or to liquidate positions. Banks may be able to change their portfolios within ten days. Some life offices may be doing very well if they even know what their liability positions were ten days ago. Claims, even for life offices, may take weeks or months to be notified, years to settle. Assurance policies are, as we have said, virtually untradeable, so we consider that a run-off approach has much to be said for it.

The opener notes that: "If we use an option pricing approach to set the basic cost of the guarantee, this has already fully costed the potential loss to the company, even if a hedge is not in place." We disagree. The pure option price is like a net premium; it makes many assumptions about the future that may not be fulfilled. Our view is that the hedging must be embedded in a real world model, to estimate the range of possible losses to the company, arising from hedging error, transaction costs, parameter uncertainty, and (in the case of GAOs) mortality uncertainty. Simple option pricing is not enough.

The opener asks whether our option pricing model was calibrated to market option prices. The answer is 'not at all', but of course it could be. He also wonders how 'fat tails' (leptokurtosis in the distributions) should be tackled. Typically, option pricing models assume Brownian motions, so that prices are available at every time instant, and when a price moves from, say, P_1 to P_2, every intermediate price is also passed through. However, in the real investment world

neither of these conditions applies; prices are only available when markets are open, and then only at short discrete intervals; prices always change at the best by small jumps ('ticks'), but may well also show much bigger jumps, which one can represent by fatter-tailed distributions in discrete time. Therefore, our approach at present would be to include fat tails in the simulated real world model (though we have not done this yet for GAOs), but to try out hedging strategies derived from models based on Brownian motions. If, in the future, hedging models based on jump processes are developed, one could use them.

Mr Clay raises the question as to how much an up-to-date actuary might have known in the mid-1980s, and others raise the question as to how quickly new ideas come into practice. We have pointed out that the MGWP reported in 1980. Its report received a lot of attention at that time. A later working party, chaired by John Ryan, carried out a survey in 1985-86 into how life offices had dealt with maturity guarantees. Unfortunately its report was never completed and never published, but one of its findings was that most life offices had used the MGWP methodology in setting up statutory reserves, and that none had found any particular difficulty in implementing it. It could be confirmed from the returns of life offices to the Department of Trade which ones had used stochastic simulation methods to calculate contingency reserves for maturity guarantees on the lines of the MGWP report and when they had done so; but our understanding is that the MGWP methodology was widely understood and used. Further, Guidance Note 8 by 1985 referred to allowance for stochastic variations (as we quote in ¶12.3); this is a little oblique, and does not mention simulation methods directly, but we would interpret it to be allowing for that possibility.

Therefore, adding to Mr Clay's comments and replying to Mr Duval and others, we suggest that the main principles of quantile reserving were established within a very few years after their major presentation in the MGWP report (and in earlier papers, such as those by Benjamin), and quite quickly became part of established actuarial and life office and supervisory practice in relation to maturity guarantees. It remains a puzzle to us as to why actuaries on all sides did not apply the same methods to GAOs at a much earlier date.

Mr Abbott comes at the matter with an accounting approach, which he understands better than we do; but our impression is that accountants do not seem yet to have got to grips with how to treat what we call 'contingency reserves'. One needs enough to cover 'technical provisions', amounts that will almost certainly be needed for policyholder liabilities. There may also be a 'free estate', belonging to the shareholders, which they can do what they like with, and perhaps take away if they choose. In between there are contingency reserves, the reserves for adverse deviation, including solvency margins, some or all of which *may* be needed to pay the policyholders' liabilities, but to the extent that they are *not* needed will revert to the shareholders. These intermediate reserves are, of course, essential if one wishes to conduct any insurance business, which is necessarily risky; but they do not clearly belong either to the policyholders or to the shareholders. How should they appear in the accounts? Should accountants not develop a third category that would include these?

There is also the question of our approach of relating the fair value to the initial premium (net of relevant expenses) (see ¶12.10). In that context we were assuming that there were enough life offices for there to be an active market at least in the selling of insurance contracts, that their prices were reasonably correctly estimated and reasonably similar, so that all participants would agree that the policyholder charge was 'fair' in that context. It is, indeed, the price at which the liability is transferred between willing parties (the policyholder and the life office). In practice, however, some offices may underestimate the costs of a contract and may unwittingly undercharge; others may intentionally offer 'loss leaders'.

Further. in practice the prices at which insurance products are sold in the market seem to vary considerably, so much so that economists would describe it as a very inefficient market. It is an interesting question how great this inefficiency is, and why it persists.

Mr Abbott questions whether what a company actually does (on the asset side, or in relation to hedging) matters in accounting terms. We disagree. The best estimate, the fair value and the prudent reserve all depend, as we have explained, on the investment policy that

the office is able to, and intends to, carry out. So, that is a matter both for accountants and for supervisors.

Mr Duval asks how new methodologies should be adopted by the profession. Mr Reynolds and Mr Wells also comment on this, and Mr Goford gives an interesting, and relevant, historical precedent. We suspect that this will be even more of an issue in the future than it has been in the past, given the greater scrutiny to which our practices are being subjected. We have no answer to this question. It may be that the Actuarial Standards Board will have a key role to play here, as several speakers suggest. However, this problem is not unique to actuaries, and it may be helpful to see how others, for example the medical profession, deal with it.

Professor Rogers suggests that expected utility would be a useful tool. We agree; but we have used it in Section 8.2, where the formula in ¶8.2.5 in fact uses a (kinked linear) utility function, although we have not used that terminology.

We quite agree with Mr Reynolds that market price is not a universal guide. It is always only a marginal price. The action of hedging, in sufficient quantity, may affect the prices of the underlying securities. We discuss this in ¶¶10.6.7 to 10.6.9. Mr Bergman seems to have experienced this problem in practice.

Mr Bergman also makes many sensible practical points. If policyholders are able to choose tax-free cash and usually do so, then one needs to allow for the GAO applying to only the 75% annuity part; but if the annuity is extremely far into the money, one might need to consider how many policyholders will observe this and give up the tax-free cash for a very good (but taxed) annuity. Trivial annuities: do policyholders have the choice to retain the annuity if it is extremely far into the money? If so, they may do so. Payment methods: typically policyholders may prefer to receive pension annuities monthly, but if the GAO is well into the money, a sufficiently large annuity payable annually in arrears may be more attractive; and the policyholder with several sufficiently large policies could stagger them so as to get one every month or quarter, each paid annually. Retirement ages: again, retirement annuities and personal pensions can (we understand) be drawn when the policyholder wishes, rather than when he/she retires, so the observant policyholder may well choose a well-into-the-money GAO at age 65 rather than earlier or later, if his circumstances allow it. Sufficient practical experience is needed to observe policyholder behaviour in all these circumstances; but the circumstances may never arise enough to find out.

Mr Wells makes two very good points about volumes; the small life office may have little expertise and little ability to get a good price for 'reinsurance'. A large life office may be too big to get a good price without moving the market. However, all new products start small; if the hedging is started when the product is designed, then excessive size should not be too great a problem except that the hedging action itself may move the market.

The closer suggests that our use of a unit-linked policy as an example leads us to conclude that hedging has problems. This is not the case at all. It is the absence of long-term (sufficiently) risk-free bonds that is the problem. A with-profits policy requires just the same hedging strategy as our unit-linked one. The differences are, first, that the 'share price' of the with-profits contract, including the sum assured, reversionary bonus and some estimate of the terminal bonus, changes more slowly than the unit price of a unit-linked policy, so rearranging the hedge portfolio because of that price change can be less frequent (but rearranging it because of interest rate changes needs to be just as frequent); and, second, that when reversionary bonus is declared, it is possible (but perhaps unfair) to include the cost of additional hedging as part of the 'cost of bonus' and to reduce the rate of bonus a little on that account. We suggest 'unfair', because we believe that the office should estimate the costs of the GAO in advance and include a charge in the premiums, rather than clawing it back as the policy develops. On the other hand, for a wholly with-profits office the profits and losses on the GAO portfolio are attributable to the policyholders who are deemed to be sharing in this account, so the amount available for bonus overall does need to take this into account first.

The close matching rules for unit-linked policies in fact do not affect the hedging, because the amount to be invested in 'shares' under the option pricing approach is always the full amount, including the value of the option (see ¶10.2.8), and this is always larger than the face value of the

policy. However, if legislative constraints prevent offices from hedging sensibly, then 'reinsurance' with a third party, or setting up a non-insurance subsidiary to act as the counterparty seems to be the way to effect the necessary 'legislatory arbitrage'.

Reference

Maccaulay, F. R. (1938). *Some theoretical problems suggested by the movements of interest rates, bond yields and stock prices in the United States since 1856*. National Bureau of Economic Research, New York. Reprinted by UMI Books on Demand, 2000.