JIA 120 (1993) 131-152

J.I.A. **120**, I, 131–152

STOCHASTIC SIMULATION IN LIFE OFFICE SOLVENCY ASSESSMENT

BY M. R. HARDY, B.SC., F.I.A.

ABSTRACT

In this paper an asset/liability model is used to compare the quality of information available from a set of stochastic simulations with a traditional deterministic sensitivity test approach.

The traditional approach applied to a range of variants of the basic model office fails to distinguish adequately very risky strategies from relatively secure strategies. The stochastic simulation method succeeds in ranking the various strategies considered into an intuitively satisfactory order of insolvency risk, as well as giving quantitative information on the relative probabilities of insolvency of different strategies and on the timing of potential solvency problems.

KEYWORDS

Life Office; Solvency; Wilkie Model

1. INTRODUCTION

WHILE most United Kingdom life insurance companies use model offices to project their emerging cash flows, few employ stochastic simulation methods. The preferred approach is to project the business deterministically and to use sensitivity tests to assess any insolvency risk. This method has been notably advocated by Brender (1988) in his work on a solvency standard for the Canadian industry, but with many insurers being forced to slim their reserves in order to remain competitive, is this deterministic, generally rather *ad hoc* approach adequate? In contrast with Brender's conclusions, the Faculty of Actuaries Solvency Working Party (1986) recommended strongly that insurers should, once techniques were suitably developed, be required to demonstrate to the supervisory authorities that their estimated probability of ruin is acceptably low, using stochastic simulation. Pentikäinen & Pesonen (1988) also prefer this approach.

Ross (1989) demonstrated the use of stochastic asset/liability models in practice. Although he mentions the possibility of using the model to investigate solvency, he does not give any examples of how this might be done.

In this paper, a model office is used to demonstrate the difference in the quality of information available from a set of stochastic simulations compared with a traditional deterministic approach. The model has been kept deliberately simple, both for practical reasons, and to avoid obscuring the results with the intricacies or idiosyncrasies of a more complex model. However, the model does allow for the essential dynamic interactions between the assets and the liabilities, and the results are sufficiently striking to be worth serious consideration. The major conclusion of the paper is that scenario testing of an office using deterministic projections may seriously underestimate the true insolvency risk. This conclusion, I believe, transcends the simplicity of the model.

The asset model used is the Wilkie investment model (Wilkie, 1986), which has recently been the subject of investigation by a joint working party of the Institute and the Faculty (Geoghegan *et al.*, 1992). In § 5.4.4 of their report, the working party state:

> "It is worth noting . . . that the model may not be considered appropriate for the estimation of extreme values, such as probabilities of ruin, etc."

The main problems with the Wilkie model, which led the working party to the conclusion that the model is inaccurate in the tails, are that the assumption of normally distributed residuals is questionable, and that the inflation part of the model is 'too tame', as it does not allow for shocks or for sustained periods of high, increasing inflation. But the tails do not have to be particularly accurate if we use them to compare relative strategies rather than to find an accurate, single probability of ruin, nor, as we shall see, do they have to be particularly accurate to give far better information on insolvency risk than is available from standard non-stochastic projections.

2. The Model Office

The model office used is a mature, ongoing office which sells 25-year with-profits endowment assurances and 25-year non-profit term assurances to lives aged 30. Bonus is awarded in the form of guaranteed reversionary bonus during the term of the policy, plus a terminal bonus on a claim by death or maturity. Bonus on previously declared bonus is always twice the rate of bonus on the basic sum assured. The terminal bonus is assessed such that participating policyholders receive 95% of their asset share on maturity, unless the guaranteed liabilities exceed this, in which case the terminal bonus will be 0%. It is assumed that the office's investments are held in U.K. gilts and U.K. equities.

The model office employs a dynamic investment strategy, which follows Ross (1989). The office is assumed to invest 80% of its assets in equities while its asset/liability ratio (A/L) exceeds 1·25. If the ratio falls below 1·25 the office moves progressively into gilts, reaching 100% in gilts if the A/L ratio is less than 1·05. The rationale for this strategy is that the higher gilt investment allows the office to use a higher rate of interest in the valuation of liabilities, reducing the liability value and thereby improving the solvency position on the A/L basis.

Reversionary bonus is also determined dynamically, by applying 20% of the excess of asset shares over the guaranteed liabilities, smoothing the resulting rates using 3-year moving averages, and rounding to the nearest $\frac{1}{4}$%. Valuation reserves are calculated on the statutory maximum interest basis, allowing for the dynamic asset allocation. New business is assumed to continue at past levels of growth for 5 years, after which the total business of the office is run off as a closed fund. This reduces the need for a premium setting strategy (I have not used one) without using the unsatisfactory approach of treating an ongoing office wholly as a closed fund when it is clearly not. A projection period of 25 years has been used, to track the progress of the office over the duration of the business written at time 0.

Six versions of the model office are considered:

(1) A base office as described above.
(2) A high equity office, in which, as long as the A/L ratio is greater than 1·25, the proportion of assets invested in equities is 90%, compared with 80% for the base office. Once the A/L ratio falls below 1·25 the office moves progressively into gilts, reaching 100% in gilts when the A/L ratio falls below 1·05.
(3) A low equity office, in which the guaranteed liabilities are matched with gilts, and the remaining assets are invested in equities.
(4) A high reversionary bonus office, in which 35% of the excess of asset shares over guaranteed liabilities is applied to fund reversionary bonus each year (compared with 20% in the base office).
(5) A high payout office, in which the total benefit paid to participating policies at maturity is the greater of 105% of asset shares and the guaranteed liabilities (compared with 95% in the base office).
(6) An office with higher equity exposure (85%), higher reversionary bonus (30% of (asset shares—value of guaranteed liabilities)), lower endowment assurance premiums (90% of the base office) and higher asset share payout (100%) than the base office. Below, this office is sometimes described as the 'high all' or the 'high everything' office.

All of these offices are assumed to start the projection period with the same assets and liabilities, apart from the high everything office (6), which is assumed to be charging lower premiums on its in-force business than the other variants. Variations in investment, terminal bonus or reversionary bonus policy are assumed to begin at the start of the projection period. The initial ratio of assets to liabilities (A/L) for all the offices is 1·87, and the ratio of assets to liabilities + statutory solvency margin + mismatching reserve (A/StL) is 1·54. Liabilities are calculated, broadly, using the statutory interest basis and 2·5% zillmerisation. It is useful to calculate the reserves at, approximately, their minimum value, since this is consistent with the rather loose use of the term 'insolvency', by which I mean that the valuation liabilities exceed the assets, at the year end. The statutory solvency margin is calculated as 4% of the valuation liabilities plus 0·3% of the total sum at risk (with appropriate adjustment to the sum at risk proportion for business nearing maturity). The mismatching reserve is calculated by applying the GAD mismatching test—that is, it is the amount sufficient to cover the additional net liabilities arising from a 25% fall in equity prices and a simultaneous 3% fall or rise in gilt yields.

It is not suggested that these offices are particularly realistic, nor that they are equally feasible. The low equity office, since the regulations allow a higher rate of interest for calculating reserves, can compete with the base office on reversionary bonus, but would be rather uncompetitive on terminal bonus and on total payout. It is included because the strategy follows traditional textbook investment principles, and the results are rather interesting.

Further information about the model and assumptions is given in the Appendix.

3. Deterministic Testing

Under deterministic projection using best estimate assumptions, all 6 versions of the office appear to remain quite comfortably solvent. In Figure 1 the A/L ratios for the first 10 years of projection are plotted for all 6 offices, showing little difference between them, with the high reversionary bonus (4) and high everything (6) offices slightly weaker than the other 4. The high equity office keeps track with the base office, as its higher investment income is balanced by the lower yield used to calculate the liabilities. The reason for the jump in the low equity office in the 1st year is the higher (maximum) valuation rate of interest allowed by the higher gilts holding.

Sensitivity tests for adverse mortality, lapse, growth, claim amounts and expenses experience, derived largely from Brender (1988), show that these offices are not very vulnerable in these areas. In Figure 2, A/L ratios for the first 10 years of projection are plotted for the base office run assuming:

(a) best estimate assumptions;
(b) mortality deteriorates at 3% p.a.;

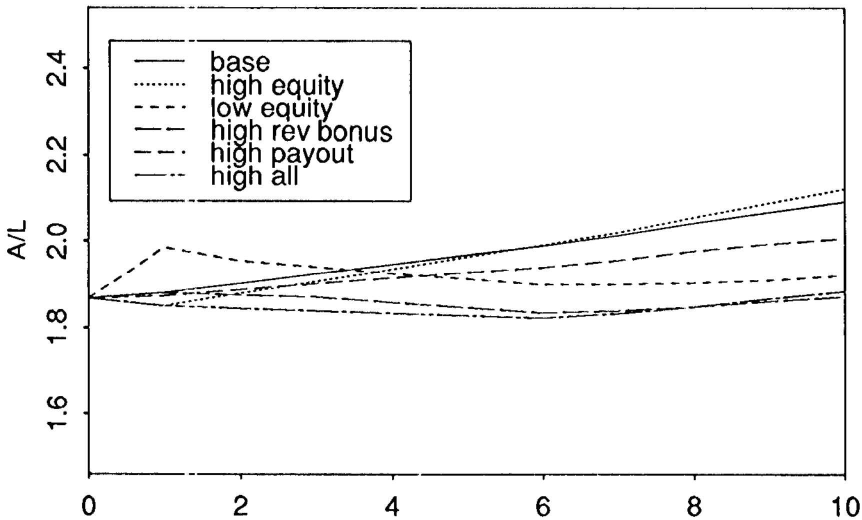

Figure 1. 10-year best estimate projections of all offices.

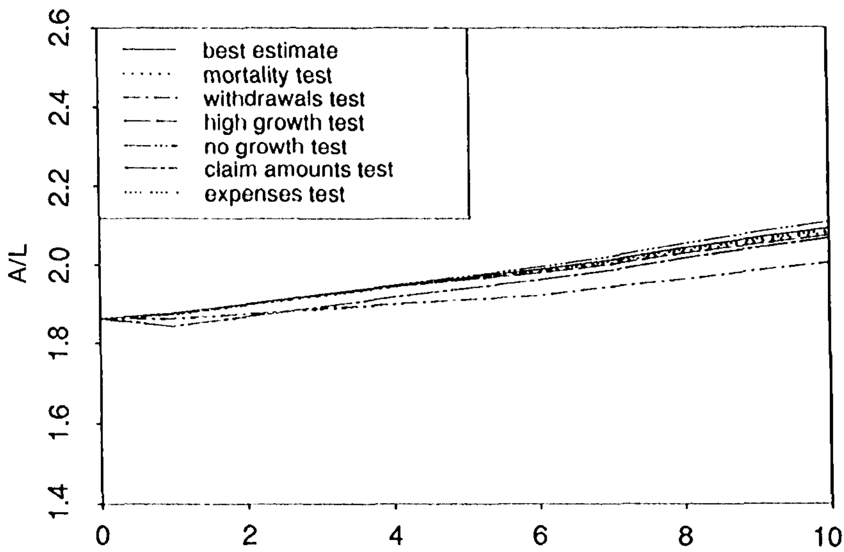

Figure 2. Base office 10-year sensitivity tests.

(c) lapse rates of half the best estimate assumptions;
(d) growth rates (in terms of policy numbers) double the best estimate assumption;
(e) growth rates of 0%;
(f) claim amounts during the first projection year of 4 times the average claim amount; this approximates to the Brender scenario test, which requires claim amounts in the 1st projection year at the upper 95th percentile of the claim distribution. In this model, subsequent basic sums assured are reduced, so that the total basic sum assured is consistent with the starting value assumed;
(g) expenses increase at inflation +3%, compared with inflation only in the best estimate run.

The offices are most vulnerable to movements in investment returns. Two tests of adverse investment experience were used. The first shows the effect of a gradual reduction in yields on both equities and gilts, resulting in a fall in the net overall yield to around 4% over the first 10 years of the projection, compared with a level net yield of 9·8% in the 'best estimate' projection. The second assumes a fall of 35% in equity prices in the 1st year of the projection, followed by a 50% recovery in the 2nd year, and a return to time 0 levels in the 3rd year. Gilts and dividends are assumed unaffected.

Figure 3 shows the effect of the first test on the 6 offices. This again shows the high reversionary bonus office and the high payout/bonus/equity office to be

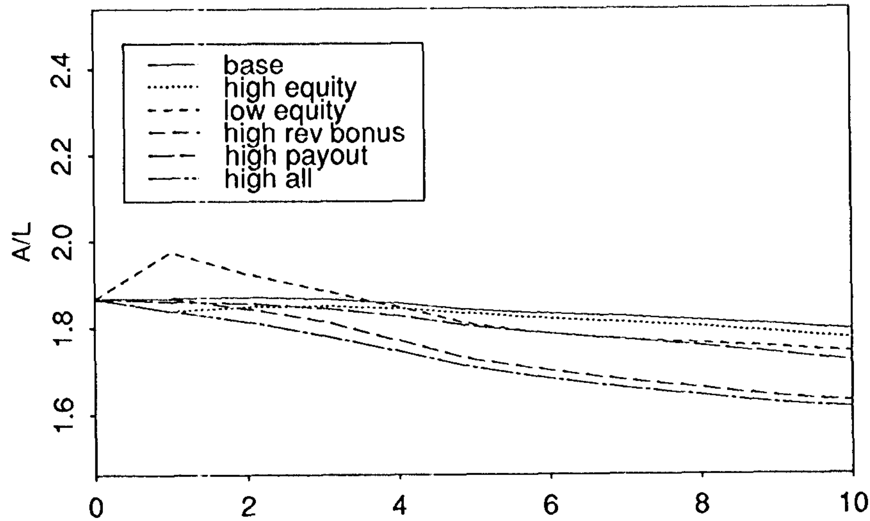

Figure 3. 10-year reducing yield projections of all offices.

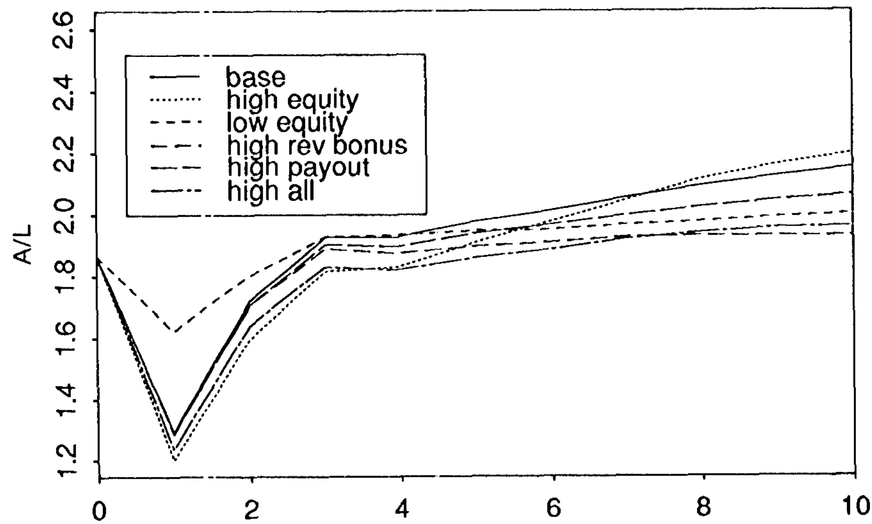

Figure 4. Equity price fall test, 10-year projections, all offices.

slightly weaker than the other 4, but after 10 years the A/L ratios of all 6 offices are still greater than 1·6, and the A/StL of all 6 offices is greater than 1·25.

Figure 4 shows the effect of the equity collapse test. In the 1st year it is clear that the A/L ratio of the low equity office is rather more comfortable than those of the other 5, which all lie between 1·2 and 1·3. Despite the rather drastic fall in equity prices, the A/StL ratios of all the offices do not fall below 1·0. By the 4th year all the offices appear to have recovered, and by the 8th year the high equity office is the most comfortable, and the low equity office is ranked 4th of the 6 variants.

These two investment tests are designed to be similar to the sort of sensitivity tests of investment return that an office might use to assess its insolvency risk. Neither test would cause any concern about the solvency risk of any of the variant offices.

4. Stochastic Simulation

In the stochastic simulation of these offices the Wilkie (1986) investment model was used for investment scenarios (including inflation rates), and individual policy claim amounts were simulated as Pareto distributed random variables. All other elements of the simulation, including claim numbers, were treated deterministically. The equity and gilt yields used in the deterministic runs were consistent with the mean yields emerging from the stochastic simulation. Each office's cash flows were projected for 25 years using the same set of investment simulations, so that the results are directly comparable. The investment strategy, reversionary and terminal bonus strategies and valuation interest rate are all calculated dynamically as described in Section 2.

In deciding the number of simulations to run, a balance must be struck between practicality and confidence in the results. If the number of simulations performed is n, then the number of simulations, X, in which the A/L ratio falls below 1·0 is binomially distributed (since each simulation either does or does not do so) and $X \sim B(n, p)$ where p is the true (unknown) probability of insolvency on the given asset/liability model. We know that $E[X] = np$ and $V[X] = np(1-p)$, so $E[X/n] = p$ and $V[X/n] = p(1-p)/n$. If we run n simulations, and find that the proportion resulting in insolvency is $\hat{p}$, then a 95% confidence interval for the true value of p for the given asset/liability model is, approximately:

$$\hat{p} \pm 1{\cdot}96 \times \sqrt{\frac{\hat{p}(1-\hat{p})}{n}}.$$

Thus, if we are working at insolvency probability levels of around 5%, we would need around 1,800 simulations for a 95% confidence interval of 1% either side of $\hat{p}$. I have used 1,000 simulations for each variant, which gives a confidence interval of 1·35% either side of $\hat{p}$ at 5%, 1·86% either side of $\hat{p}$ at 10%, and 2·5% either side of $\hat{p}$ at 20%. However, since exactly the same 1,000 investment scenarios were used for each of the variants, the effects of sampling error are very much reduced. (I have heard it suggested that 50 simulations will give an adequate indication of insolvency probability. At the 6% level—3 insolvent runs—the 95% confidence interval is (0·0%, 12·6%), which ranges from the implausibly solvent to the unacceptably risky.)

The proportions of simulations in which the A/L ratio or the A/StL (Assets/Statutory Liabilities) fell below 1·0 at one or more year end over a 25-year projection for the 6 offices are given below. An office with an A/L ratio of less than 1·0, I have called insolvent. An office with an A/StL ratio of less than 1·0 is, in principle, in some difficulties with the supervisory authorities.

Office	Insolvent simulations A/L < 1·0 (%)	A/StL < 1·0 (%)
(1) Base office	7·1	33·4
(2) High equity office	13·6	41·9
(3) Low equity office	0·1	3·6
(4) High reversionary bonus office	15·0	54·3
(5) High payout office	10·8	43·8
(6) High equity/payout/reversionary bonus, low premium office	20·7	60·0

The spread of results between the variants here is very much greater than indicated by the deterministic tests, and the ordering is much more intuitively satisfactory. Also, these numbers indicate that some of the variants are unacceptably risky. The 21% insolvency probability of the high everything office is clearly unsatisfactory, and so are the 14–15% insolvency probabilities of the high equity and high reversionary bonus offices, despite the dynamic strategies which stop the high equity investment and high reversionary bonuses when the office appears to be in trouble. None of the deterministic tests gave any indication of these levels of risk, or indeed of any insolvency risk at all. The A/L figures are perhaps a better indicator of solvency risk than the A/StL figures. The statutory liabilities (StL) are defined as the sum of the valuation reserves, the statutory solvency margin and the mismatching reserve calculated by applying the GAD mismatching test. Some—but not all—of the problem simulations arise from equity crashes, when there is reason to believe that the mismatching requirements would be relaxed, so the apparent probability of supervisor intervention is overstated here. It is unlikely that the statutory minimum valuation basis would be weakened in these circumstances. Nevertheless, it is interesting that the high reversionary bonus office is in considerably more danger of alerting the supervisors than, for example, the high equity investment office, although both have similar solvency problems on the A/L basis.

In addition to the simple insolvency probabilities, stochastic simulation gives information on the distribution of the projected cash flows and asset/liability ratios, and on the timing of potential solvency problems.

Figures 5(a) and 5(b) show the spread of projected A/L ratios for the base office. Figure 5(a) plots 50 of the 1,000 simulations, to give an illustration of the

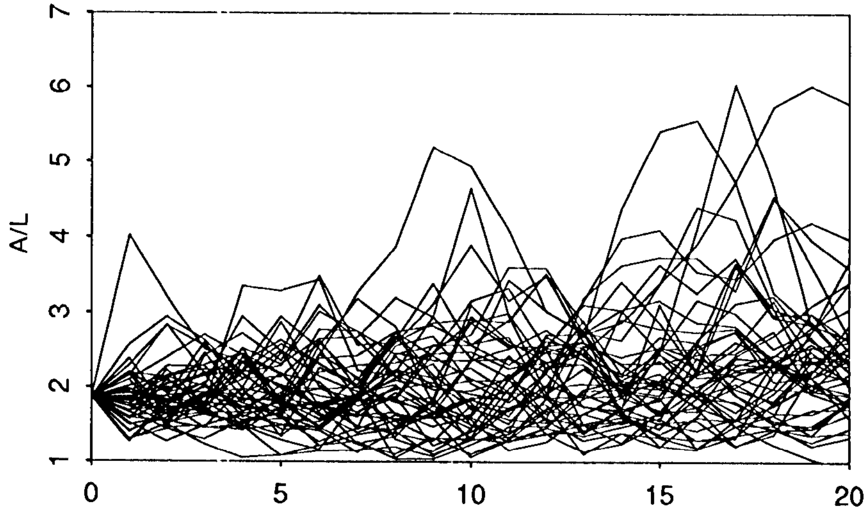

Figure 5(a). 50 Simulations of the base office.

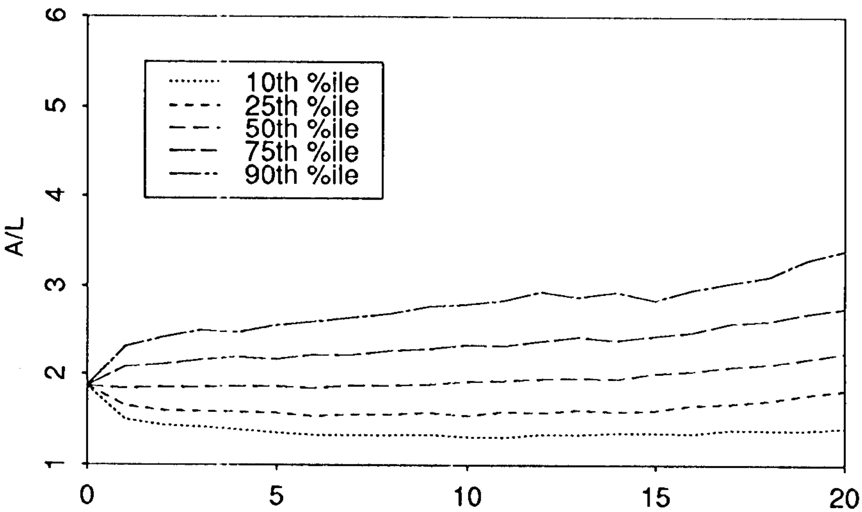

Figure 5(b). Percentiles of base office simulations.

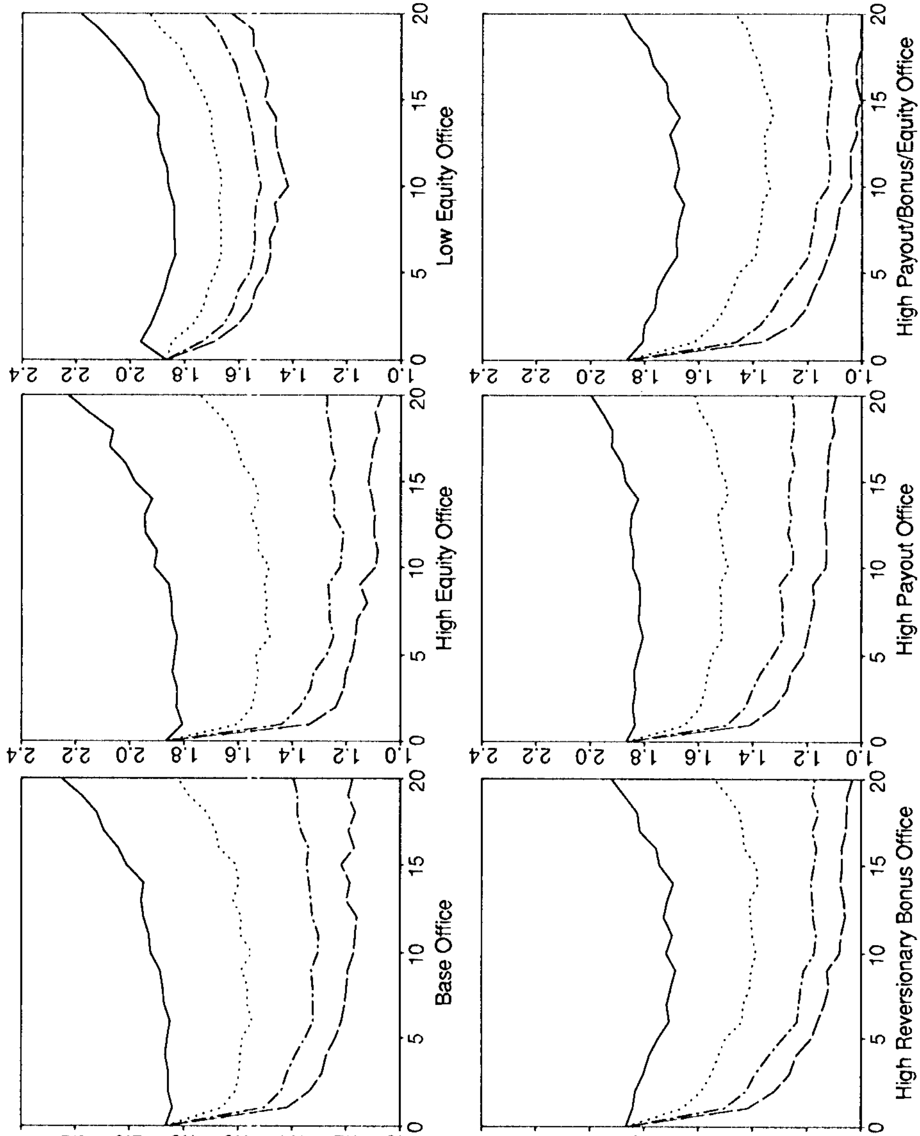

Figure 6. 5th, 10th, 25th and 50th percentiles of projected A/L ratios, all offices.

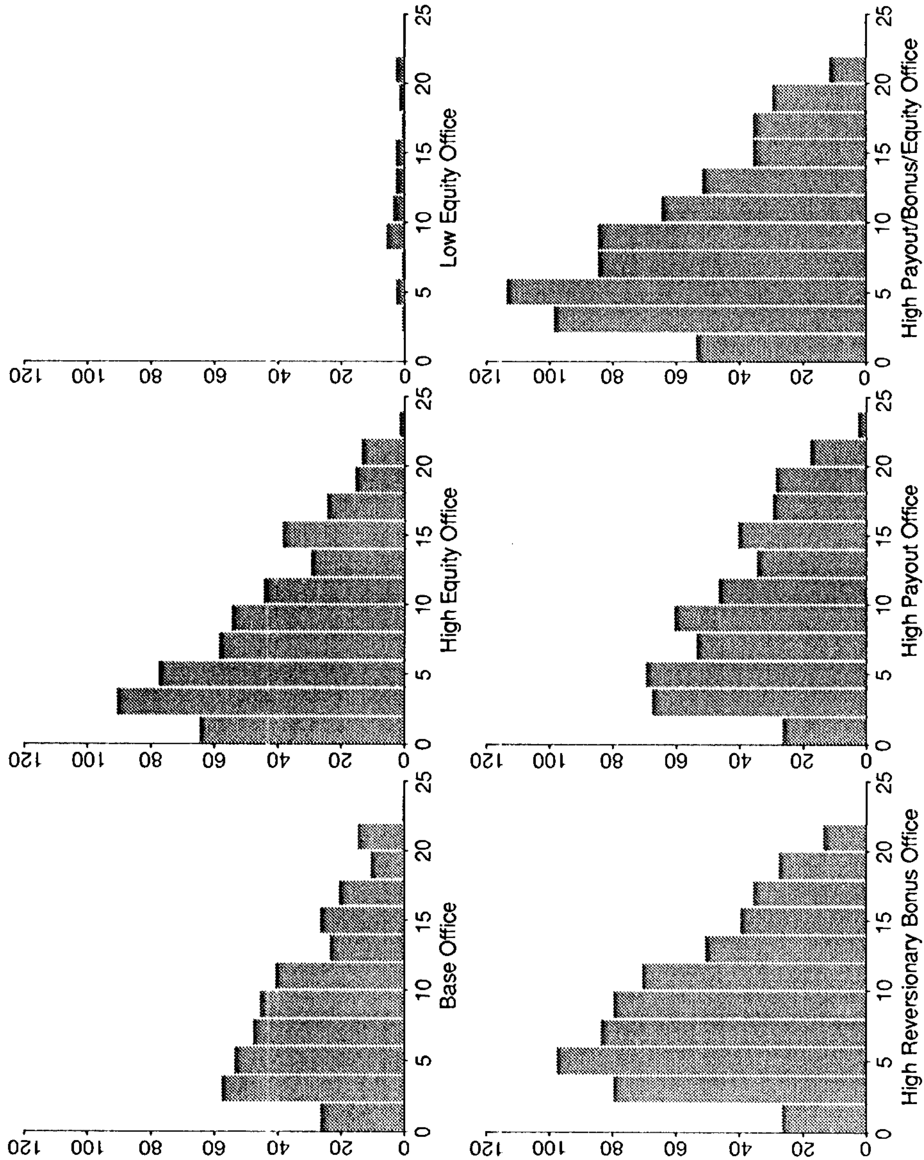

Figure 7. Histograms of years in which A/L < 1·25 for the first time.

projections produced by the model. Figure 5(b) plots, for each projection year, the median of all 1,000 projected A/L ratios, together with the 10th, 25th, 75th and 90th percentiles. This demonstrates that the median A/L ratio is fairly steady, with a slight upward trend after 15 years. The lower 10th percentile projection, based on A/L ratios, stays level at around 1·2, while the upper increases gradually over the period of the projection.

As we are concerned with solvency, the upper percentiles are not very useful. In Figure 6 the 5th, 10th, 25th and 50th percentiles of all 6 offices are shown for comparison. The comparatively small range of outcomes for the low equity office is very apparent. The high equity office median A/L ratio does not appear to vary very much from the base office over the projection period, but the lower percentiles are clearly lower. The high reversionary bonus office shows a smaller spread from the 5th percentile A/L ratio to the median than the base office, but the higher guaranteed liabilities have depressed the whole range. The high payout office has a median A/L curve that rises more slowly than the base office, and the lower 5th percentile is lower. In the high everything office the median has the 'U' shape of the high reversionary bonus office, and the lower percentiles fall fairly quickly to the insolvency level.

Figure 7 gives histograms for all the offices of the year in which the A/L ratio first falls below 1·25. The base office and the high payout office have similar shapes, but the high payout office has slightly more offices in trouble at each stage. The high equity office has very similar frequencies of low A/L ratios to the high payout office after the 8th year of the projection, but significantly higher numbers up to the 8th year. The extra risk of the high equity strategy over the high payout strategy apparently affects the early projection years most strongly. The problems in the high reversionary bonus strategy arise slightly later, on average, than the high equity strategy, although overall numbers are similar. The histogram for the high payout/bonus equity, low premium office has a similar shape to the high reversionary bonus office, but with a greater frequency of low A/Ls in the first 6 years.

5. Investment Scenarios Leading to Low A/L Ratios

The introduction of investment scenarios generated by the Wilkie model has produced dramatically different results to those obtained from the deterministic tests (the stochastic treatment of claim amounts has only a very small effect on the results). In this section we will consider what makes a dangerous investment scenario, and whether this information can be used to construct a more informative sensitivity test for use in deterministic testing, to replace the subjective tests of Section 3.

Figure 8 shows the pattern of inflation rates, net rates of return on equities and net rates of return on gilts (Consols). The means of these rates are plotted using the full 1,000 simulations used in the stochastic model runs, and the 100 simulations which gave the lowest minimum A/L ratios for the base office, that is

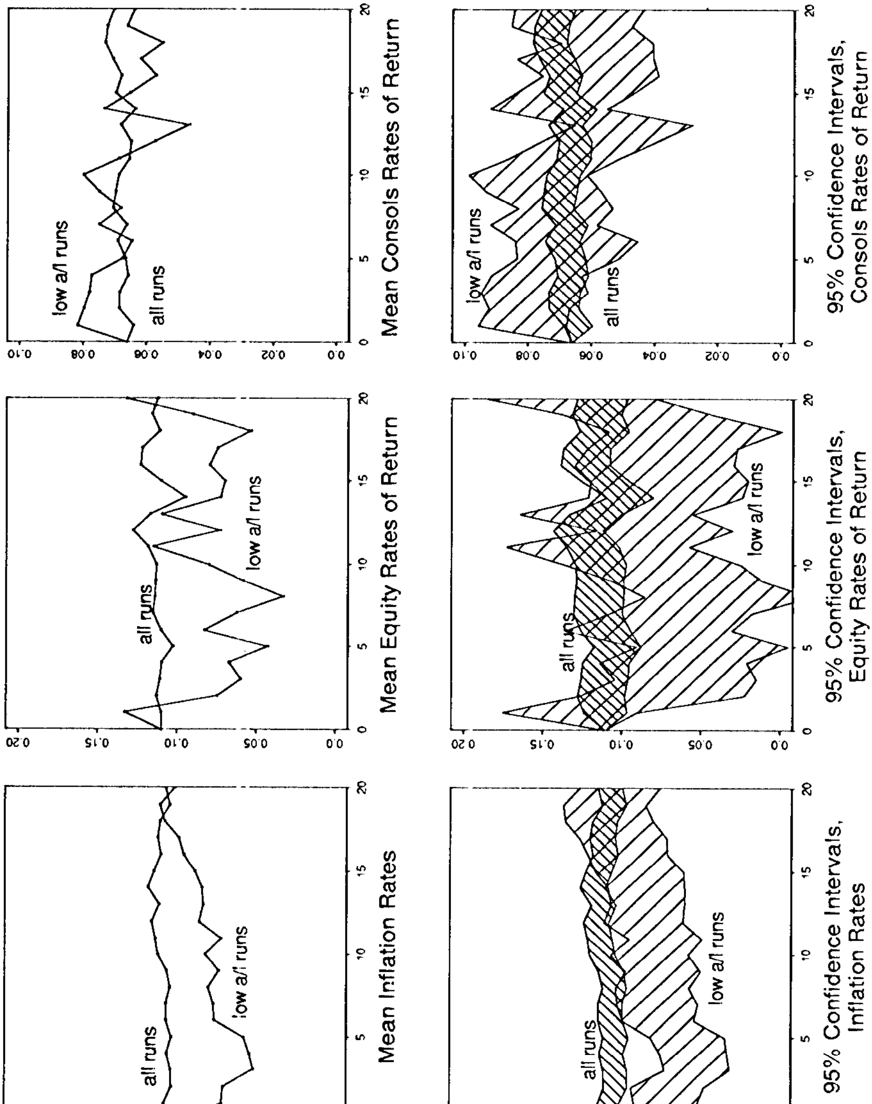

Figure 8. Comparison of investment simulations, low A/L and all runs.

the 100 simulations in which the A/L of the base office falls at some stage to less than 1·04. (These 100 investment scenarios are referred to below as the base office low A/L scenarios.) Also, for each of the projected rates, 95% confidence intervals are plotted for the mean from the full set of simulations and the mean from the low A/L set. These show that the problem lies with the runs in which the rate of inflation in the first 15 years is low, and in particular when the rate of inflation falls sharply in the first 5 years. The effect of this fall is, initially, a rise in equity returns, followed by a sharp fall, to roughly half the overall mean rate of return on equities. The effect of the initial drop in inflation rates on the Consols yield is to generate higher than average rates for the first few years, after which there appears to be no difference in the mean rates of the full set of simulations and the mean rates of the low A/L set. These early negative correlations of inflation and equity returns, and inflation and Consols returns are discussed in Section 5 of Wilkie (1986).

Figure 8 is perhaps, rather surprising. Despite the fact that expenses are linked to inflation, and that after the 5th year there is no new business, it is low inflation, not high inflation, that is a problem, because of the association between inflation and equity returns. Also, although at low A/L levels the yield on Consols is increasingly important (because of the Ross investment strategy), there is no evidence here of an association between very low A/L ratios and low Consols yields.

It has been suggested that, instead of having to run a large number of simulations, it should be possible to encapsulate the 'interesting' runs in a single, or a few, non-stochastic simulations. Geoghegan *et al.* (1992), for example, in §D.3.15 suggest,

"A better approach would be to develop a technique which condenses the results of a larger number of simulations into a relatively small number of 'representative' scenarios, each of which would be assigned a weighting factor."

Can we use the information on low A/L investment scenarios to construct a deterministic test of solvency? We could test the variants deterministically using the medians of the base office low A/L inflation rates, equity price index and dividend yield index. The results of such a test are shown in Figure 9, and demonstrate that this test is no better at distinguishing the higher and lower risk offices than were the subjective tests of Section 3.

The test does give, up to the 15th year, roughly the same ordering of risks as the full stochastic runs, with the low equity office being the least at risk, followed by the base office. The variant with the highest risk is the high everything office, with the remaining 3 variants being fairly closely bunched between that and the base office. However, the test does not give low A/L values in any years, for any variant. Since at least one of the variants has a very high risk of insolvency, the test is clearly inadequate. One reason why this test does not result in any particularly low A/Ls is that the danger in many of the base office low A/L runs arises, not from the mean equity returns of those runs, although these are lower

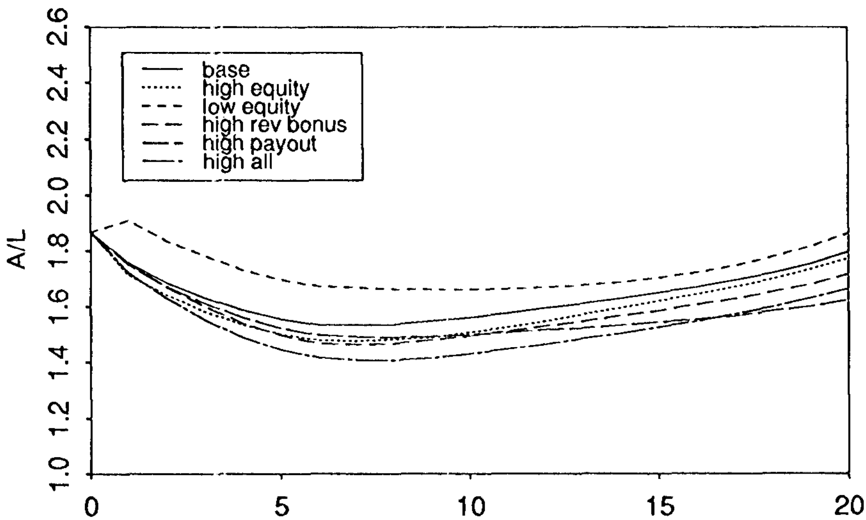

Figure 9. Deterministic test using low A/L parameters.

than other runs, but from the relatively higher frequency in these runs of dramatic falls in equity yields. If we define an equity crash to be a sudden drop in equity prices of more than 25%, there were, in the full set of 1,000 simulations, an average of 1·6 crashes per 25-year simulation, and in the 100 base office low A/L simulations, an average of 2·5 crashes per 25-year simulation.

However, even if we construct a deterministic test with the characteristic peaks and troughs of the stochastic runs, it will not differentiate very effectively between the low risk and high risk offices—no single scenario can give apparently adequate information, even on the relative severity of the insolvency risks of the 6 variants, and the absolute levels of the risks are quite obscured. This information emerges only with the accumulation of results from the full 1,000 scenarios tested.

Some examples of how the projected A/L ratios of the offices vary under individual investment scenarios (chosen randomly from the base office low A/L scenarios) are shown in Figure 10 to demonstrate this point. None of these, individually, yields any more information than the subjective, deterministic tests of Section 3; but aggregated, the 1,000 simulations give the far more useful information of Section 4.

Since the low points of the A/L curves do not, in general, retain the ordering of risks of the full set of simulations, it would not even be possible to obtain reliable results in comparing these variants using all the 100 base office low A/L runs, since these are only defined to be dangerous for the base office, and in many cases

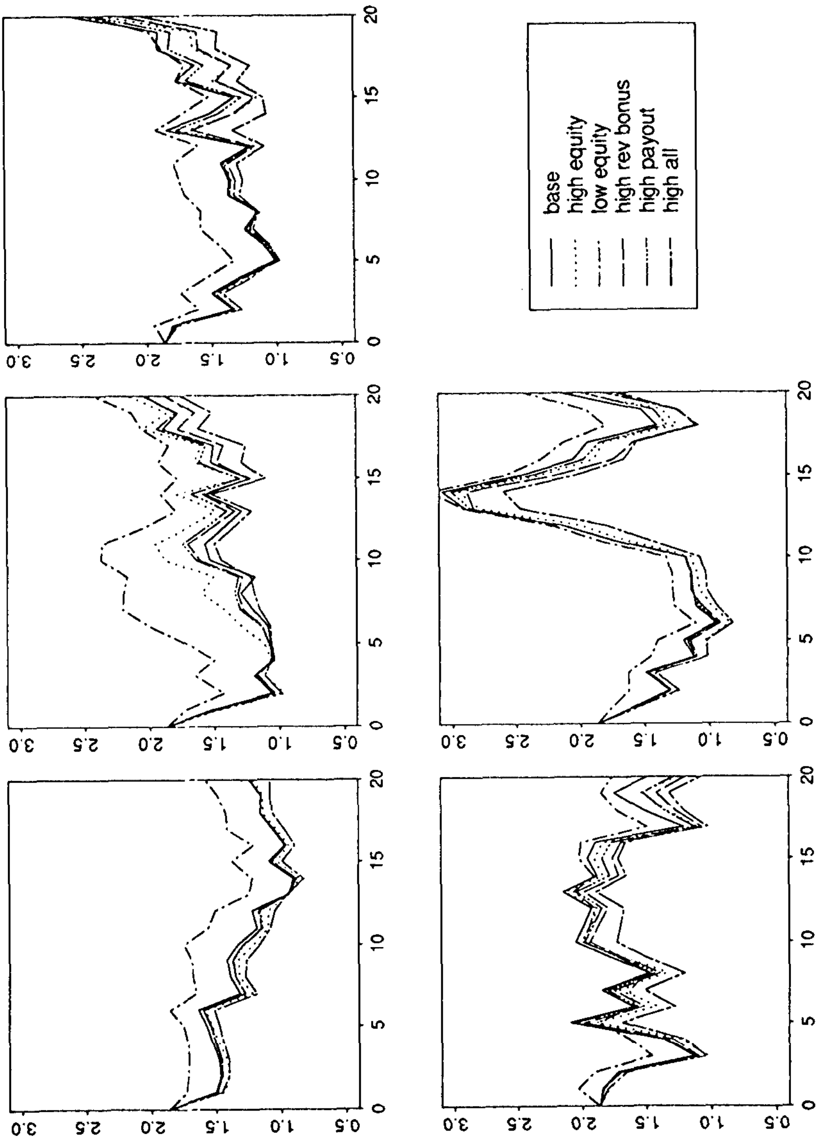

Figure 10. Examples of individual simulations, all offices.

different scenarios are dangerous for different variants. Considering only the 100 base office low A/L runs, we have the following numbers of insolvent simulations for the 6 variants:

Office	Insolvent simulations (from the 100 worst base office runs)
(1) Base office	71
(2) High equity office	89
(3) Low equity office	1
(4) High reversionary bonus office	80
(5) High payout office	85
(6) High equity payout/reversionary bonus, low premium office	89

This has not preserved even the basic ordering of risks of the full set of simulations, and has given no indication of the absolute levels of risk generated by the full set of simulations.

Also, because the dangerous investment scenarios differ for the same office with a different initial estate, we cannot even assume that the dangerous scenarios for the base office in 1 year could be assumed to be the 100 most dangerous scenarios in the following year. Thus, I see no way of circumventing a full set of stochastic simulations by using a few hand-picked scenarios to define a deterministic test.

6. Conclusions

Stochastic simulation is not very different from the scenario testing proposed by Brender, and used extensively in U.K. life offices. In place of 1 or 2 prescribed investment scenarios, we use here 1,000 scenarios. The prescribed scenarios are supposed to represent some kind of worst-case situation, and are not assigned any probability. An office either passes or fails each test. The stochastic scenarios are supposed to represent a wide range of future investment conditions, and no attempt is made to identify dangerous scenarios in advance. The only assumption is that each scenario is equally probable.

For this simple model, the deterministic tests gave the strong impression that none of the variants considered was under any risk of insolvency at all. Yet, for the highest risk office, the probability of insolvency before the business in force has run off is estimated at 21%, unacceptably high by any standards. For these variants the traditional tests could lead to a misleading assessment of the true level of risk, both in absolute terms and in relation to other offices.

For these conclusions to have any meaning in practice, we need to be persuaded that the deliberate simplicity of this model office does not invalidate it. Although the range of policies is very restricted, and allowance for some factors (notably tax) is very approximate, the model has the essential features of more complex models. The dynamic strategies cover the main influences on solvency—investment strategy, bonus declaration and the valuation rate of interest. They are not supposed to be optimal in any sense, but they are supposed to represent reasonable, credible reactions to future situations.

It is also necessary to accept that the Wilkie model generates a set of feasible potential investment scenarios. This is being considered elsewhere, but I believe the results of the comparison of deterministic and stochastic testing are so striking that it is not necessary to assume that the model is absolutely accurate, only that the range of results, especially considering the lower 20% tail, is broadly credible.

Arguments against stochastic testing appear frequently in actuarial writing. The argument that it obviates actuarial judgement (see Geoghegan *et al.* [1992], § D.2.15) is no more true of this type of model than any other actuarial model, from the simplest life table to the most complex office projection model. Like all models, stochastic simulation provides a tool—a very powerful tool—to assist the actuary in her or his analysis. Actuarial judgement is essential, both in the choices of asset and liability parameters for the model, and in interpretation and presentation of results.

Actuarial judgement will certainly be needed also in the original construction of the asset/liability model. A common argument against stochastic simulation is the computer run time required. This argument is often based on the premise that the asset model is to be slipped into an existing large, comprehensive, CPU draining liability model, and that the only way to keep the run time within practical limits is to keep the number of simulations very small. Perhaps a better route would be to use a simplified liability model for stochastic simulation, retaining the existing model for more detailed deterministic work. Some research would be required into a suitable range of sample policies, and the sensitivity of the model to the simplification of, for example, tax or mismatching reserve calculations. It is possible that a simplified liability model, which allows an adequately large number of simulations, could be more accurate and more flexible in application than a highly complex liability model run on a very small number of simulations. The model used in this paper takes around 120–200 minutes to run 1,000 simulations on a SUN SPARCstation.

Geoghegan *et al.* (1992) suggest that stochastic simulation involves using 'horrendously incomprehensible techniques to prove something blindingly obvious'. I do not believe that the results of Section 4 are 'blindingly obvious', and I hope that, as the techniques become more widely used, it will become clear that they are very far from 'horrendously incomprehensible'.

References

ABBOTT, W. M. (1992). Statutory Regulation of Long Term Business. Monograph, Institute of Actuaries and Faculty of Actuaries.

BRENDER, A. (1988). Testing the Solvency of Life Insurers. *Proceedings of the 2nd International Conference on Insurance Solvency.*

FACULTY OF ACTUARIES SOLVENCY WORKING PARTY (1986). The Solvency of Life Insurance Companies, *T.F.A.* **39**, 251.

GEOGHEGAN, T. J. *et al.* (1992). Report on the Wilkie Stochastic Investment Model, *J.I.A.* **119**, 173.

PENTIKÄINEN, T. & PESONEN, M. (1988). Stochastic Dynamic Analysis of Life Insurance. *Proceedings of the 23rd International Congress of Actuaries*, **Volume 1**, 421.

ROSS, M. D. (1989). Modelling a With-Profits Life Office. *J.I.A.* **116**, 691.

WILKIE, A. D. (1986). A Stochastic Investment Model for Actuarial Use. *T.F.A.* **39**, 341.

APPENDIX

THE MODEL OFFICE AND PROJECTION ASSUMPTIONS

The office is assumed to write 2 types of policy, with-profits endowment assurances and non-profit term assurances. In the deterministic runs, sums assured of the endowment policies, at time 0 values, are assumed to be £14,657, and sums assured of the term assurances are assumed to be £29,314. In subsequent years these are assumed to increase at the inflation rate. These are the means of the truncated Pareto distributions used to simulate claims on and off in the stochastic simulations. The endowment assurance sums assured, in the stochastic runs, are assumed to have Pareto distribution with parameters (in the 1st year) $\alpha = 2{\cdot}5$ and $\lambda_0 = 22{,}500$, truncated at £300,000, adjusted for inflation for entrants in years other than the first. The term assurance sums assured in year 1 are assumed to have Pareto distribution with parameters $\alpha = 2{\cdot}5$, $\lambda_0 = 45{,}000$, truncated at £600,000. The initial distribution function for the sums assured of policies entering at time t is then:

$$F(x) = \begin{cases} 1 - (\lambda_t/(\lambda_t + x))^{\alpha} & 0 < x < A \\ 1 & x \geqslant A \end{cases}$$

where λ_t is λ_0 increased in proportion to the inflation index, and A is the truncation point.

At the start of each year in which new business is written, the total sum assured for the new cohort is simulated, using this distribution as the basis for the individual policy claim amount—the total for the large claim numbers involved here is actually simulated using a normal approximation to the sum of individual truncated Pareto distributions. This total is then reduced in subsequent years by the simulated sums assured of the exits, and the distributions of the remaining sums assured are adjusted (by adjusting the λ parameter) such that the total sum assured of the stayers plus exits is the same as the total sum assured of the starters. In this way if, for example, large sums assured are simulated in the early years of a projection, they will be compensated in later years with lower expected sums assured, so that the overall total is consistent with the starting point.

New entrants are assumed to continue for the first 5 years of the projections, assuming a rate of growth in policy numbers of 3%, and in sums assured (E[sums assured] in the stochastic runs) of the appropriate inflation rate.

At time 0 of the projection, 5,000 new endowment policies and 5,000 new term assurance policies are assumed to be written. In other years the number of new policies written is assumed to be 10,000 $(1{\cdot}03)^t$ with 50% in each policy type, where $-24 < t \leqslant 4$.

All policyholders are assumed to enter at age 30, and policy terms are all 25 years.

Premiums are £40·86 per mille for endowment policies and £2·50 per mille for term assurances.

Withdrawal rates assumed in all projections are, for endowment assurance policyholders, 10% in the 1st year, 8% in the 2nd, 5% in the 3rd, followed by 2% in every year up to the 24th, with no withdrawals in the final year. For term assurance policyholders withdrawal rates of 4% are assumed throughout. It would be interesting to link lapses with investment conditions in the stochastic runs, but, in the absence of any data, this would be too speculative to be helpful.

On surrender of a with-profits policy, the office is assumed to pay out 90% of the policy asset share. On surrender of a non-profit policy no surrender value is paid.

On a claim by death or maturity, it is assumed that the office pays terminal bonus sufficient to make up the sum assured plus attaching bonus to 95% of the asset share (approximately for claims by death, exactly for claims at maturity). If the guaranteed liabilities exceed the asset share, no terminal bonus is paid. The reversionary bonus declarations in the past are assumed to have been constant at 2·5% of sum assured plus 5% of previously declared bonus. In all the projections, bonus on bonus is maintained at twice the level of bonus on the sum assured. The method of calculation is described in Section 2. On a deterministic projection of the base office (using 'best estimates') the reversionary bonus is stable at 2·5%/5·0% for 5 years, increasing gradually to 4%/8% after 15 years. The terminal bonus on maturing policies, which is assumed to be 55% at the start of the projection, is projected to increase to 70% by the 10th year, falling back to 50% after 20 years, as the business runs off.

Mortality of 95% of A1967–70, and lapses at the same rates as those for time 0 given above, have been assumed in creating the time 0 in-force business.

Mortality of 95% A1967–70 is assumed throughout the projection period.

Expenses at renewal are 2% of premiums, plus 0·0525% of sums assured. The 2nd term increases in line with inflation (deterministic or simulated, as appropriate). Initial expenses for endowment policies are 3% of the sum assured plus 30% of the initial premium, and for term assurances are 80% of the initial premium. Apart from the inflation link, expenses are treated deterministically in the stochastic runs.

The Wilkie model, run on the reduced standard basis (Wilkie, 1986), with a minimum inflation rate of −5%, has been used to generate the investment conditions for the stochastic simulations—that is, where the inflation rate generated by the unadjusted model is less than −5% a rate of −5% is assumed. For deterministic projections a net yield of 6·4% on gilts and 10·6% on equities is assumed throughout. These are consistent with the mean rates of return assumed in the stochastic runs.

The ratio of assets to asset shares at the start of the projection is 1·35.

Currently minimum reserves and solvency margins are governed by the

Insurance Company Regulations 1981. A concise description of these is given in Abbott (1992). The valuation reserves in this model are calculated using a rate of interest of 92·5% of the yield earned on assets over the previous year, allowing redemption yields on gilts and dividend yields on equities, subject to a maximum of 7·2% gross of tax. This is consistent with the 1981 regulations.