Incorporating individual life company variation in simulated equity returns

Mary R Hardy
Department of Actuarial Mathematics and Statistics
Heriot-Watt University
Edinburgh EH14 4AS
U.K.
Telephone 44-31-451-3210
Fax 44-31-451-3249
E-mail mary@cara.ma.HW.ac.uk

Summary
In applying the Wilkie investment model within a stochastic asset-liability model, it is generally assumed that the rates of return generated satisfactorily model the rates of return earned by individual life companies. This limits the possibility of simulating the performance of an individual company relative to other companies, since the return on investments, particularly equity investments, may be the major source of difference between, for example, a solvent life insurer and an insolvent life insurer. In this paper the equity rates of return generated by the Wilkie model are interpreted as the rates of return on the equity market as a whole. The performance of an individual company's equity portfolio is assumed to be related to, but not identical to the market performance. After consideration of a small set of data on life insurers' equity returns over the past 8 years, a model for the relationship between the company equity yield and the market equity yield is proposed. An illustration of the effect on solvency, using a simple model life office, is then presented.

Prise en compte des variations entre les compagnies d'assurance vie dans les simulations des rendements boursiers

Mary R. Hardy
Département de mathématiques et de statistiques actuarielles
Heriot-Watt University
Edimbourg EH14 4AS
Royaume-Uni
Téléphone 44-31-451-3210
Télécopie 44-31-451-3249
E-mail mary @cara.ma.HW.ac.uk

Résumé

Lorsque que l'on applique le modèle d'investissement de Wilkie dans le cadre d'un modèle stochastique de gestion actif-passif, on suppose généralement que les taux de rendement obtenus correspondent de manière satisfaisante aux taux de rendement obtenus par les compagnies d'assurance vie individuelles. Ceci limite la possibilité de simuler les performances d'une compagnie donnée par rapport aux autres, étant donné que le rendement de l'investissement, en particulier les placements en actions ordinaires, peut être la principale source de différences par exemple, entre une compagnie d'assurance vie solvable et une compagnie d'assurance vie insolvable. Dans la présente étude, les taux de rendement boursiers générés par le modèle de Wilkie sont interprétés en tant que taux d'intérêt du marché boursier dans son ensemble. On adopte l'hypothèse que le portefeuille de titres de la compagnie considérée est lié mais pas identique aux performances du marché. Sur la base d'un petit ensemble de données relatives au rendement du portefeuille boursier de compagnies d'assurance vie au cours des huit dernières années, nous proposons un modèle permettant de rendre compte de la relation entre le rendement des titres de la société et le rendement du marché. Nous présenterons ensuite une illustration de l'effet sur la solvabilité en utilisant un modèle simple de compagnie d'assurance vie.

1 Introduction

The Wilkie investment model [Wilkie (1986)] is used widely to model investment returns by researchers who use stochastic simulation to investigate life office financial control (see for example [Ross (1989)]). In this paper we consider the effects of incorporating in a stochastic model life office a slightly different interpretation of the equity returns simulated by the Wilkie model. Here we are interested particularly in asessing the solvency of the model life office.

The equity returns are particularly important in the solvency context, since UK life insurance companies which write substantial volumes of conventional business may be rather more vulnerable to adverse variation in the equity returns than to adverse movements in inflation rates or in rates of return on fixed interest securities. This is because high inflation is associated with, and can be more than compensated by, high rates of return from equity investments, and because the amount of variability of returns from fixed interest securities is small compared with that of equities. A demonstration of this is given in [Hardy (1993)].

The Wilkie model was designed to simulate dividends and prices of the FT-All Share Equity Index, and parameterized (broadly) using past data from this index. To what extent does this index represent the equity investments of an individual UK life office? In interpreting these simulated rates as the returns earned by an individual office we are in difficulty if we want to compare offices. If we use the Wilkie model rates within a stochastic simulation which attempts to compare one office with another, or an individual company with the life office market as a whole, we must either use the same investment simulations for each office - implying that each office is assumed to achieve exactly the same rates of return on their equity and fixed interest portfolios - or we generate different sets of simulated investment conditions, implying that offices are operating in different markets. What is required is to be able to model the different rates of return earned by individual offices, within a consistent market simulation framework. By consistency we mean that, for example, it is unlikely that one office will experience an equity crash, and another office an equity boom in the same projection year, but it *is* likely that, if a crash is simulated for the market as a whole, one company's equities will be affected to a different extent than another's. The differences may be purely random, or may have a systematic element, as a result of different investment strategies.

2 Life Office Equity Performance

Information on the rates of return achieved by UK life offices on the equities held in respect of conventional business does not exist in reliable and accessible form. It is easy to obtain rates of return achieved on the linked funds of the life offices, since these are published in the UK trade press; the drawback is that, comparing these published data with some sample Department of Trade returns, it is clear that different companies treat provision for tax on capital gains differently, which makes comparison between companies invalid.

Capital gains tax does not provide problems if we consider the published returns on linked, tax-exempt UK equity funds (available for example from the *Pensions Stats* pages of Pensions Management magazine). It is not clear to what extent these rates represent the returns achieved by the same companies on their non-linked equity assets backing conventional business liabilities. Because of this uncertainty, any conclusions we draw from these data will be fairly tentative.

In figure 1 the rates of return on UK equity linked exempt funds of 30 companies are shown. The rates are annual rates, from 1 June to 1 June each year, allowing for reinvestment of income. The companies are all UK companies writing conventional business; the data cover the period 1 June 1985 to 1 June 1993. This is quite an interesting period, including as it does 2 years of the mid-80's bull market, followed by the 1987 crash.

The rates of return achieved by the FT-All Share index over the same period, allowing for reinvested income, are shown in Figure 2.

The performance of the companies relative to the index, may be investigated by looking at $R(i,j) = (1 + I(i,j))/(1 + Ift(j))$ where $I(i,j)$ is the rate of return for the ith office in the jth year, and $Ift(j)$ is the rate of return on the FT All-Share Index in the jth year, allowing for reinvestment of income.

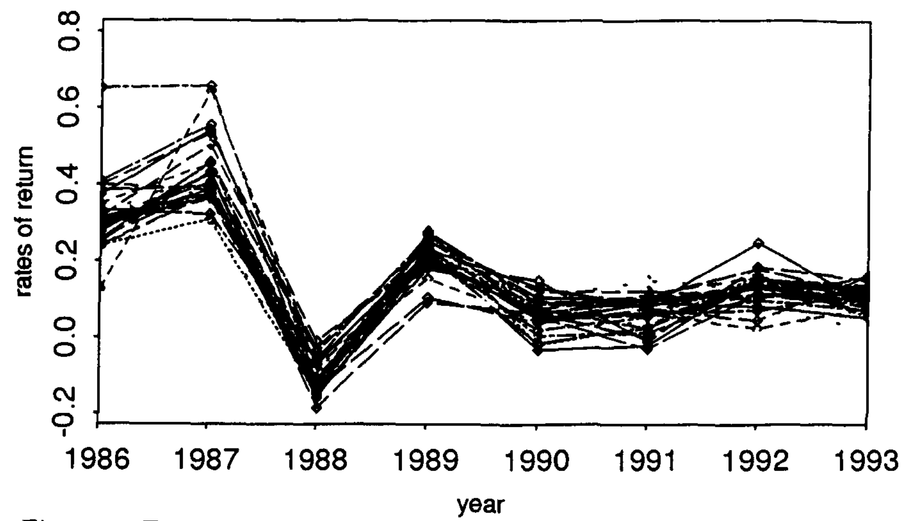

Figure 1: Equity rates of return, income reinvested, for 30 Life Offices

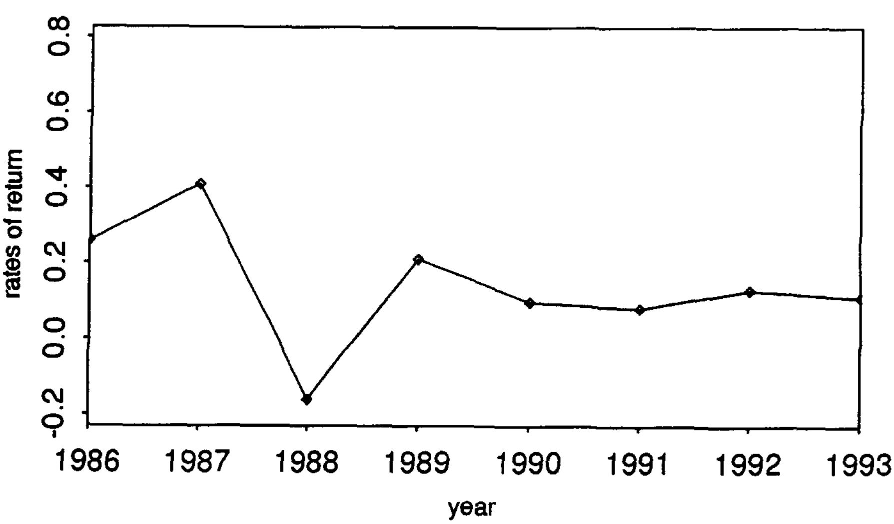

Figure 2: Rates of return on FT-All Share Index, income reinvested

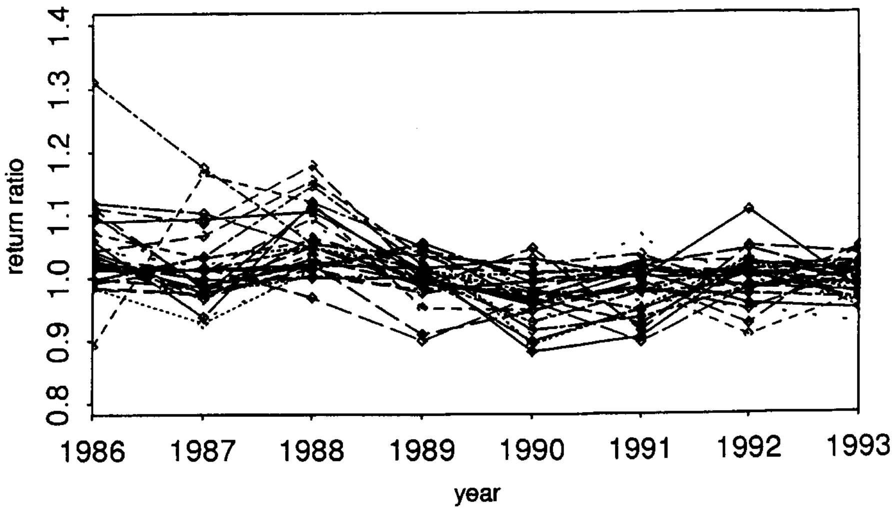

Figure 3: Ratio of offices' equity returns to FT-All Share returns

The factors $R(i,j)$ are referred to below as return ratios. In Figure 3 these ratios are plotted for the period 1985/6 to 1992/3.

In Figure 4 the return ratios data is summarised. The mean and standard deviation of $R(i,j)$ by company are given in the first plot. Including the outlier, there is a weak positive correlation between the mean and the standard deviation, with correlation coefficient 0.46. Excluding the outlier the correlation coefficient is 0.21, which does not provide convincing evidence that high risk is rewarded by high returns.

In the second plot the relationship between fund size and standard deviation of returns is given. The fund size here is the size (in £000's) of the exempt linked equity fund on which the rates used have been earned. The figure plotted is the mean of log(fund size) over the period, intended to give an indication only of the relationship between fund size and variability of returns. The correlation is around -0.5. The correlation between the standard deviation of the returns and the company size, as measured by the 1990 long term assets (actually log(long term assets)) is weaker, at -0.3.

In the third and fourth diagrams of Figure 4 the mean and standard deviation of $R(i,j)$ are plotted for the 8 years. The 1988 figure for the mean value of $R(i,j)$ shows that the companies, on the whole, managed to survive the 1987 crash better than the FT-All Share Index did. The standard deviation of the offices in the 1987/88 crash year does not look out of line, but it appears that since the crash the companies have been tracking each other more closely, at the expense of very high returns. The increasing use of derivatives may have contributed to this.

Can we use this data set to suggest how individual company equity returns may be simulated in a stochastic projection, given that we intend to treat the Wilkie simulated equity returns as the market index rate of return for the projection year?

In Figure 5 the companies have been grouped by the average size of the linked funds used, and all the 8 years' data for the companies in each group have been aggregated. Some grouping is necessary to have sufficient data for assessing the distribution of the $R(i,j)$, and this choice is an attempt to retain some homogeneity with reasonable sized data sets. All four histograms are rather too fat-tailed for a normal distribution, although the symmetric shape indicates a distibution of the same family. The logistic distribution has the symmetry of the normal distribution, but with the fatter tails required. In fact the log-logistic distribution, which in this region is very close to a logistic distribution provides a slightly better fit. It also has the advantage that it is distributed on the interval $(0,\infty)$, rather than logistic distribution range of $(-\infty,\infty)$. The curves superimposed on the 4 histograms are log-logistic curves.

The fit in each case is not bad, having a χ^2 value of between 3.8 and 6.1, with 3 degrees of freedom. The log-logistic distribution also gives a reasonable fit for other aggregations of the data. While the evidence does not overwhelmingly prove that the log-logistic distribution provides a good model for the ratio of company returns to market index returns, it does at least give some support for the idea.

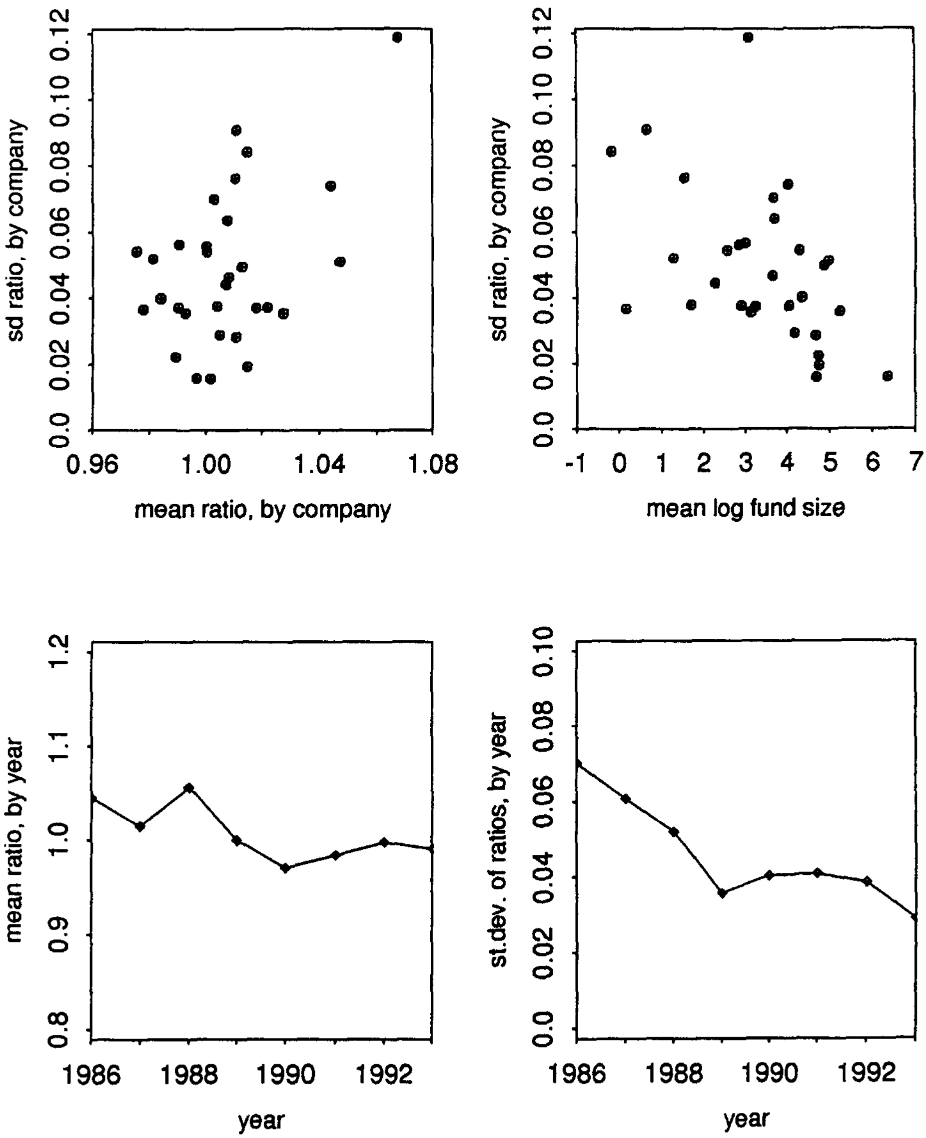

Figure 4: mean and standard deviation plots of $R(i,j)$

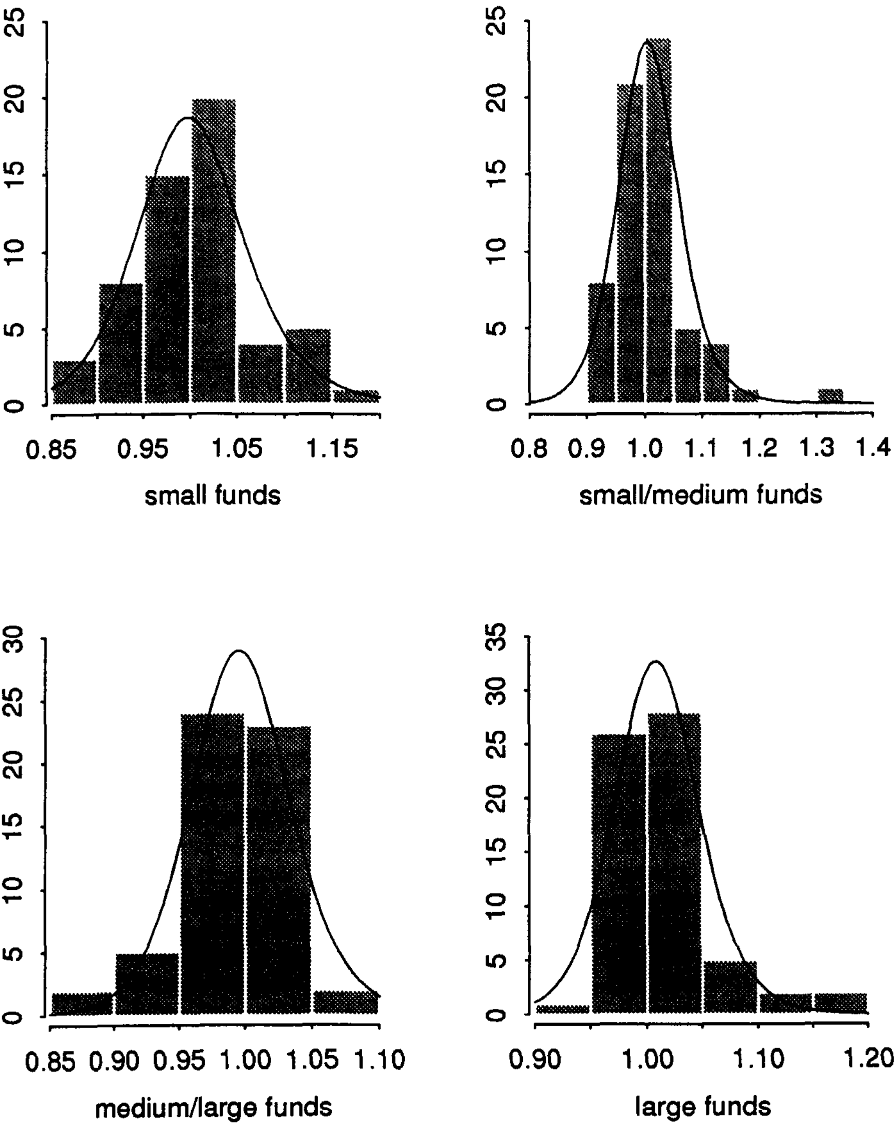

Figure 5: Histograms of return ratios, with fitted log-logistic curves

3 The Asset Model

The asset model proposed for the investment performance of an individual company employs the usual Wilkie model generated rates of inflation, gilt yields and equity price and dividend indices. The equity indices are interpreted as the market indices, rather than individual company rates. They are used to generate an annual rolled-up *market* rate of return on equities in the tth year of I_t^w, say.

The individual company's earned rate of return on equities, I_t^c is related to the market (Wilkie) rate as

$$\log(1 + I_t^c) = \log(1 + I_t^w) + \varepsilon_t$$

where $\varepsilon_t \sim logistic(\mu, \tau)$
or, in other words,

$$I_t^c = L_t.(1 + I_t^w) - 1.0$$

where $L_t \sim log - logistic(\mu, \tau)$.

The terms ε_t (and hence also L_t) are assumed independent and identically distributed, and independent of I_t^w. This means that we assume that ε_t does not depend on t – that is, we may model a *company* effect, by varying μ and τ, but we do not model a *year* effect. In fact, analysis of variance of the data in §2 indicates that there may be a year effect. We might expect the ε_t to be serially correlated, or to be dependent on I_t^w, or both. In the data of §2, there is some evidence that the ε_t may be expected to behave differently in years in which the equity market crashes than in other years. However, this is not being modelled here.

We know approximately, from simulation, the moments of I_t^w; after the first few years of each projection these are reasonably stable. The 1,000 simulations of the model used in the examples below give $E[I_t^w] \approx 0.12$ and $V[I_t^w] \approx 0.625$, implying a standard deviation of the rolled-up rate of approximately 0.25.

The relationship between the moments of I_t^w and I_t^c are easily found, given the independence assumption for L_t and I_t^w:

$$E[I_t^c] = E[L_t].E[1 + I_t^w] - 1.0$$

and

$$V[I_t^c] = V[I_t^w].E[L_t^2] + V[L_t].E[1 + I_t^w]^2$$

We know that

$$E[L_t] = \frac{e^{\mu}\pi\tau}{\sin(\pi\tau)} \text{ and } V[L_t] = \frac{e^{2\mu}2\pi\tau}{\sin(2\pi\tau)} - \left(\frac{e^{\mu}\pi\tau}{\sin(\pi\tau)}\right)^2$$

Possible values of μ and τ underlying the data in §2 are estimated by maximum likelihood, using the same aggregation of offices by fund size as in §2. We use data on the return ratios collected over the 8 years for the estimates - those in Figure 6 are based on the 8 ratios for each office, those in the table below are based on 7 or 8 offices data (7 small offices, 8 small/medium offices, 7 medium large offices, 8 large offices) aggregated. This is possible because we are assuming ε_t are i.i.d..

	μ	τ
small funds	0.000	0.036
small/medium funds	0.007	0.032
med/large funds	-0.004	0.022
large funds	0.009	0.022

The range of estimates of μ and τ over individual companies is illustrated in Figure 6. In only one company (where $\hat{\mu} = .04, \hat{\tau} = .023$) is $\hat{\mu}$ significantly different from 0, at 5% significance.

Returning to the moments of I_t^c, we use 4 example values of (μ, τ) to demonstrate the expected effect, assuming E$[I_t^w] = 0.120$ and sd$[I_t^w] = 0.250$

(μ, τ)	E$[L_t]$	sd$[L_t]$	E$[I_t^c]$	sd$[I_t^c]$	% increase in sd
(-.02,.03)	0.982	0.054	0.101	0.253	1.2%
(0,.022)	1.001	0.040	0.121	0.254	1.2%
(0,.035)	1.002	0.064	0.122	0.261	4.4%
(.02,.03)	1.0217	0.056	0.144	0.263	5.2%

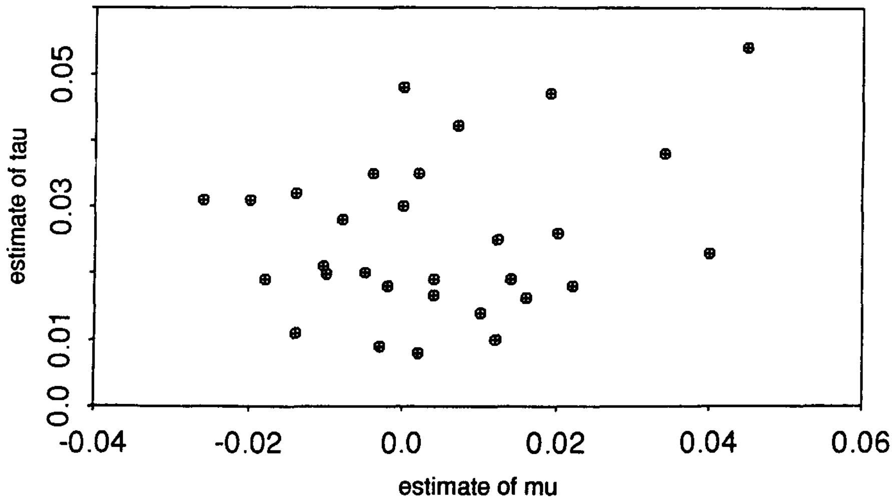

Figure 6: Individual company estimates of μ and τ

The effect of the company specific variability, for most companies, on the standard deviation of rolled-up equity returns is fairly small, an increase of about 3% over the standard deviation of the market as a whole in each year. This is consistent with the data of §2 which showed rates of return for the companies, over the 8 years, on the whole slightly higher than the standard deviation of the FT All-Share index rates. The mean standard deviation of the company data is around 5% higher than the index over the period 1985/6 to 1992/3.

To demonstrate the relationship between the simulated market rates of return and the simulated company rates of return on equities, in Figure 7 two simulations of the equity returns are shown, The four broken lines represent the returns simulated for individual companies with the given μ and τ values, the unbroken line shows the simulated market rates of return.

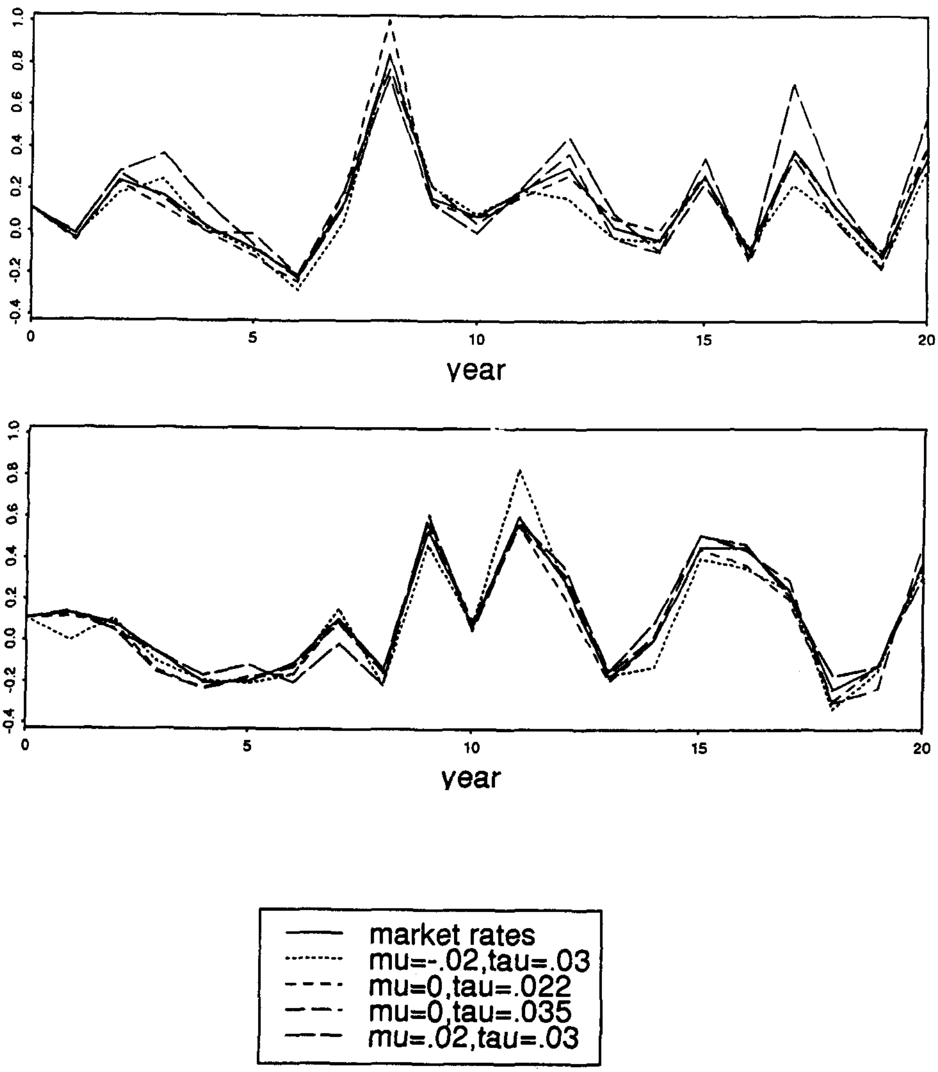

Figure 7: Two example simulations of equity returns, market and individual company

4 A Stochastic Model Office

The effect of incorporating this log-logistic adjustment into a stochastic projection of a model office will be demonstrated with a simple model office, as used (and described in more detail) in [Hardy (1993)]. This models a mature life insurer writing participating endowment assurances and non-participating term assurances to lives age 35; all policies have a 25 year term.

Reversionary bonuses are determined dynamically, by applying a proportion (17%) of the difference between the policy asset shares and the guaranteed liabilities to fund a two-tier bonus, under which the bonus on the sum assured is always half the bonus on bonus. Terminal bonuses are declared such that the payout on maturity is the greater of 95% of the asset shares and 100% of the guaranteed liabilities.

The valuation rate of interest is also determined dynamically; valuation reserves are calculated on the statutory maximum interest basis with 2.5% zillmerisation - which makes tham close to the statutory minimum valuation reserve under current UK regulations.

The final dynamic element is the proportion of assets invested in equities; the office employs a dynamic investment strategy which is derived from that of [Ross (1989)], under which if the A/L ratio falls below 1.25, assets are progressively transferred into gilts, with 100% gilt investment if the A/L ratio falls below 1.05. As long as the A/L ratio is greater than 1.25, 80% of the assets are assumed invested in equities, the remainder assumed invested in gilts. This is an important strategy, since we do not assume below any difference between offices in gilt yields. This means that, as an office falls into difficulties the equity yields become increasingly unimportant; nevertheless the allowance for individual company variation in equity returns described in §3 may have a large effect on insolvency probability.

The sums assured of new policies and of exits are stochastically simulated. The sums assured in force are assumed to have a truncated Pareto distribution, the mean of which is annually updated to allow for the simulated sums assured of exits. The sums assured of maturing policies is not stochastic, but is fixed such that, together with the exits, the total sum assured is consistent with the simulated total on entry.

Mortality is treated deterministically.

New entrants are assumed to enter for the first 5 projection years only; surrendering with-profit policyholders are assumed to receive 90% of their asset shares on withdrawal.

The office, at the start of the projection period of 25 years, has an Asset/Asset-shares (A/AS) ratio of 1.35, and an Asset/Liabilities (A/L) ratio of 1.87. The ratio of assets to the "statutory liabilities", (comprising the valuation liabilities plus the statutory solvency margin plus the mismatching reserve required to be held by UK companies by the supervisors), A/StL is 1.54.

5 Some simulations of the stochastic office

We consider 5 versions of the model office. The liabilities at the start of the projection for each variant are identical. The first office is assumed to earn on its equity portfolio the rates of return generated by the Wilkie model, without adjustment. The other 4 model offices earn L_t times the rate of return generated by the Wilkie model- that is those assumed earned by the first office- where L_t is independently simulated each year from a log-logistic distribution, with parameters μ and τ. The parameters are the same as those illustrated in §3. In Table 1 the results of 1,000 simulations, each projecting assets and liabilities for 25 years, are summarised. The offices are (in a loose sense) insolvent when A/L is less than 1.0, since the liabilities are close to the statutory minimum; they are potentially in trouble when A/StL is less than 1.0, when, although assets may cover the valuation liabilities, they are insufficient to cover also the statutory minimum solvency margin and the required mismatching reserve. These offices all operate in the same generated markets - that is, the rates assumed achieved by the first office are the assumed underlying market rates of the other offices in each simulation.

		$L \sim log - logistic(\mu, \tau)$			
	office 1 L=1	office 2 $\mu = -0.02$ $\tau = .03$	office 3 $\mu = 0.0$ $\tau = .022$	office 4 $\mu = 0.0$ $\tau = .035$	office 5 $\mu = 0.02$ $\tau = .03$
$A/L < 1.0$	6.8%	21.5%	10.3%	14.0%	4.8%
$A/StL < 1.0$	33.4%	61.3%	41.0%	46.0%	28.8%

Table 1: % of simulations insolvent or in trouble, from 1,000 simulations.

We see that the incorporation of individual office variability has a significant effect on the insolvency probabilities. Office 4, for example, has values of μ and τ that appear fairly standard from Figure 6, but the insolvency probability has doubled, moving from respectably safe to fairly high risk.

Apart from this demonstration of the range of insolvency probabilities for different companies in what we could call a "Wilkie market", a major benefit of being able to model individual company variation is that we can assess the risk that the company falls out of line with the market, in terms of payouts to policyholders, or of asset-liability ratios. In practical terms, in the UK, this may be a more realistic definition of insolvency, since it is unlikely that the insurance supervisors would allow a company which lagged significantly in either of these to continue trading.

References

[Hardy (1993)] Hardy M R (1993) Stochastic Simulation in Life Office Solvency Assessment, *Journal of the Institute of Actuaries* **120**

[Ross (1989)] Ross M D (1989) Modelling a With-Profits Life Office, *Journal of the Institute of Actuaries* **116** 691.

[Wilkie (1986)] Wilkie A D (1986) A Stochastic Investment Model for Actuarial Use, *Transactions of the Faculty of Actuaries* **39**,341.