Stochastic investement models and their actuarial applications

Autor(en): **Wilkie, A.D.**

Objekttyp: **Article**

Zeitschrift: **Mitteilungen / Schweizerische Vereinigung der Versicherungsmathematiker = Bulletin / Association Suisse des Actuaires = Bulletin / Swiss Association of Actuaries**

Band (Jahr): **- (1992)**

Heft 2

PDF erstellt am: **10.08.2024**

Persistenter Link: https://doi.org/10.5169/seals-967256

B. Wissenschaftliche Mitteilungen

A.D. WILKIE, Reigate

Stochastic Investment Models and their Actuarial Applications[1]

It was your Honorary President, Professor Hans Bühlmann, who coined the phrase "actuaries of the third kind" to describe those actuaries who are interested not just in life insurance or general insurance, but in the application of actuarial methods to modelling investments – the assets of an insurance company or a pension fund, or indeed of any other financial institution. Modern financial economists, particularly in North America, have applied very elaborate mathematics and statistics to modelling investments, but often their interests have been directed towards how the market operates, towards economic equilibrium models, and towards the investment of funds in a generalised situation, rather than the investment of funds of a particular institution, whose liabilities – like insurance policies or pension fund liabilities – are not themselves tradeable on the market.

It is up to actuaries to consider the implications of the models derived by modern financial economists to actuarial problems, and in my talk today I want to introduce the idea that stochastic investment models can be of use to the actuary of a life assurance company, whether or not they are of use to the investment manager.

I intend first to describe a very simple stochastic model, what some English speakers might describe as a "Mickey Mouse" model. I do this in order to give you the flavour of stochastic modelling without getting too complicated, or too realistic.

I shall then discuss various ways in which stochastic modelling can be of use to a life insurance company. My examples will be taken from the British situation, but I hope that you will be able to translate them into Swiss circumstances.

In the third part of my talk I shall describe in outline the stochastic investment model I have developed for use in Britain, although it is also applicable in other countries outside Britain. This is what has become known in Britain as the "Wilkie stochastic investment model".

[1] Vorgetragen an der Mitgliederversammlung der Schweizerischen Vereinigung der Versicherungsmathematiker vom 12. September 1992 in Winterthur.

1 A Simple Example of a Stochastic Model

In my simple model I want to consider only the accumulation with interest of a payment of one per annum for 20 years. I simplify, as compared with a real life assurance policy, by missing out mortality and expenses, in order to get over the principal concepts.
You are obviously familiar with the concept of an accumulation with interest of one per year, payable at the beginning of each year for n years, and accumulated at a fixed rate of interest, i. This is given the actuarial symbol $\ddot{s}_{\overline{n}|}$.
We can calculate this in a number of ways, but the one I want to use is the recursive calculation:

$$\ddot{s}_{\overline{0}|} = 0$$
$$\ddot{s}_{\overline{n}|} = (\ddot{s}_{\overline{n-1}|} + 1)(1 + i)$$

I have used a fixed value of n of 20 years. $\ddot{s}_{\overline{20}|}$ can be presented as a function of i; it increases with i, and Figure 1 shows how it increases. It goes from $-4\,\%$ to $+12\,\%$ interest, because I shall find this useful later. Of course at $0\,\%$, $\ddot{s}_{\overline{20}|} = 20$.

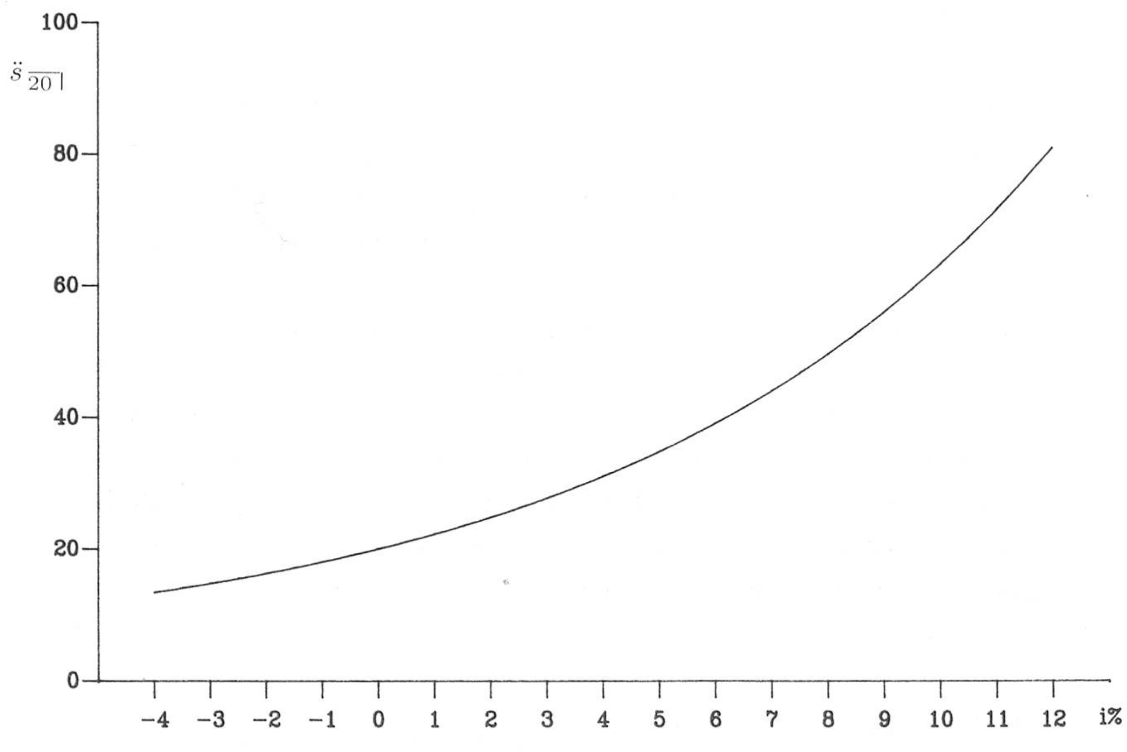

Figure 1: $\ddot{s}_{\overline{20}|}$ as a function of i

Now let us consider what happens if the rate of interest each year is not a constant i, but instead varies each year. Let us say that the rate in year t is $i(t)$. I shall treat each $i(t)$ as a random variable, and the sequence $i(1), i(2), \ldots, i(20)$ as a stochastic process, i varying with t. I shall postulate different stochastic models for the behaviour of $i(t)$, but for any particular realisation of $i(t)$, for $t = 1, 2, \ldots, 20$, we can define the accumulation with interest one per year, which I shall call $S(20)$ this time, and calculate it recursively in the same way as before:

$$S(0) = 0$$
$$S(t) = (S(t-1) + 1)(1 + i(t)) \qquad \text{for} \quad t = 1, 2, \ldots, 20\,.$$

For each particular model for $i(t)$ I am interested in the distribution of $S(20)$.
I shall first assume that successive values of $i(t)$ are independent. I shall assume that $1 + i(t)$ is lognormally distributed. To be precise I shall assume that $\delta(t) = \log(1 + i(t))$ is generated using the formula:

$$\delta(t) = \mu + \sigma \cdot z(t)$$

where the value of $\mu = \log(1.04)$, ie the median rate of interest each year is 4 %, the values of $z(t)$ are all normally distributed with zero mean and unit standard deviation, and are independent and identically distributed, and in my first example I put the value of $\sigma = 0.01$, representing roughly 1 % a year in the rate of interest. It is possible to calculate the moments of $S(20)$ analytically, but I do not know how to calculate the complete distribution except by random simulation, so I have used this method. In each case I have used 10,000 independent random simulations.
The results are shown in Figure 2. Do not worry about what α is at this stage; its value here is 0. I show the distribution in terms of the equivalent uniform rate of interest, by inverting the function $\ddot{s}_{\overline{n}|}$ at a fixed rate of interest which I showed earlier. 95 % of all simulations produced values of $S(20)$ between 29.27 and 32.78, equivalent to accumulations at a rate of interest between 3.5 % and 4.5 % over the 20 years. The remaining 5 % were split almost equally between the cells representing 3 % to 3.5 % accumulation and 4.5 % to 5 % accumulation. Only 3 out of the 10,000 simulations fell outside the range of 3 % to 5 %.
If you believed that this was an appropriate model for the investments of a life company, it would be reasonable for you to calculate the sum assured for a premium of 1 per annum on the basis of an accumulation at 3 % interest, and to allow for the expected excess interest to be distributed in the form of bonus in some appropriate way.

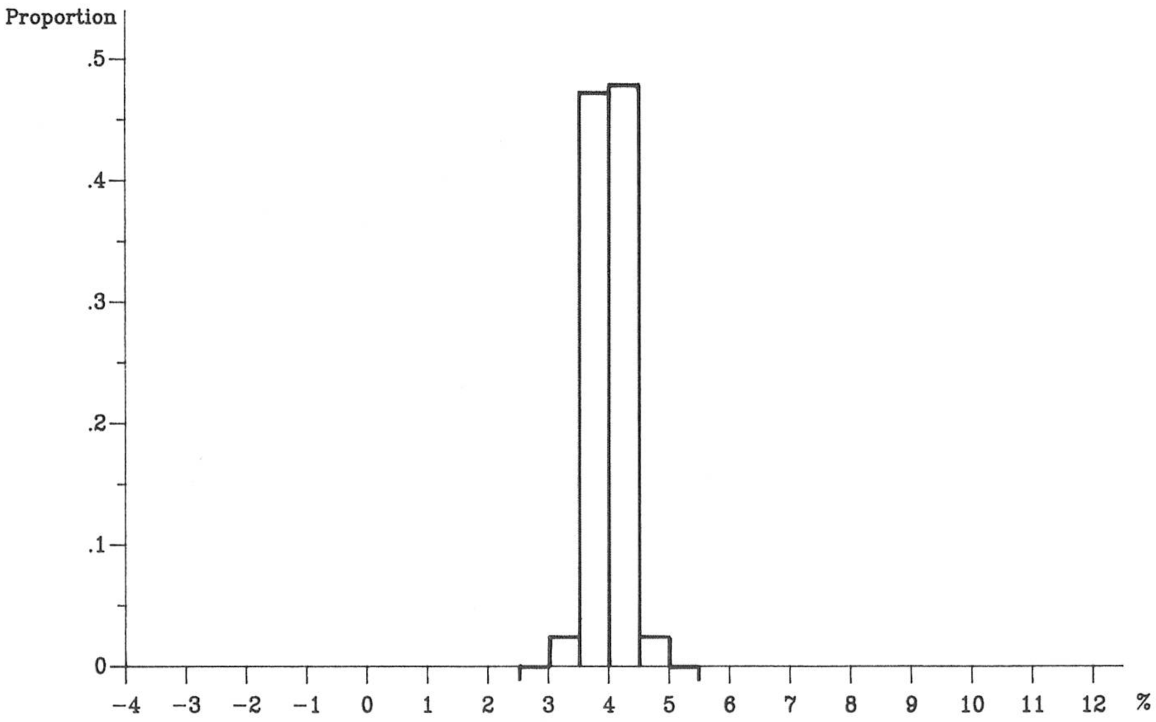

Figure 2: Distribution of $S(20)$; $\alpha = 0.0$, $\sigma = 0.01$

What happens if we change the values of the parameters. Obviously if I change the value of the median, the whole distribution moves up or down along the scale, and I shall not bother to show this. But if I double the standard deviation, σ, to 0.02, the range of results moves out quite a bit, as shown in Figure 3. Almost all the simulated values of $S(20)$ lie in the range equivalent to a rate of interest of 2.5 % to 5 %, but there are 17 out of 10,000 below that range and 16 above it. Two of these fall into the 1.5 % to 2 % range and another two into the 6 % to 6.5 % range. I would prefer to calculate premiums for a with profit contract in order to guarantee no more than 2.5 %, and again to use an appropriate bonus system for distributing the excess interest.

Now what happens if the rates of interest in the successive years are not independent but are correlated? The simplest model is of first order autocorrelation. I now put:

$$\delta(t) = \mu + \alpha(\delta(t-1) - \mu) + \sigma \cdot z(t)$$

where everything is defined as before (with the value of σ back to 0.01) except that the value of $\delta(t)$ is connected with the value in the preceding year, $\delta(t-1)$,

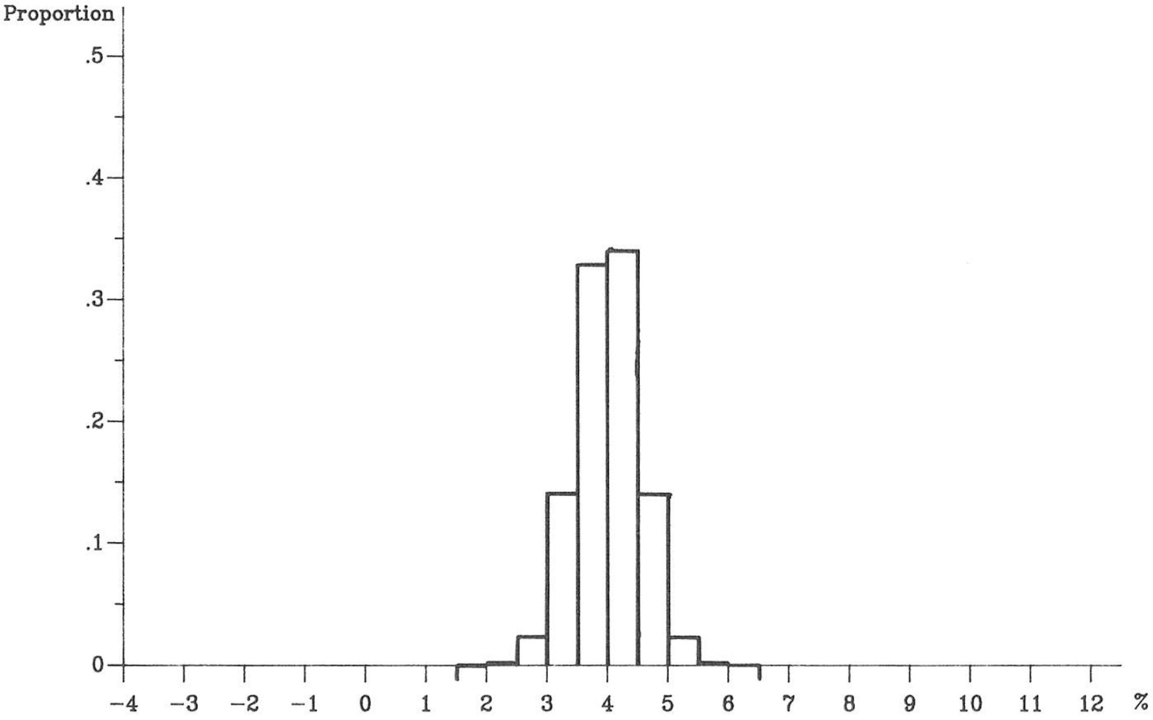

Figure 3: Distribution of $S(20)$; $\alpha = 0.0$, $\sigma = 0.02$

through the middle term on the right hand side of the equation. This term contains the α that was zero in my first two examples.

Figure 4 shows the results, using values of σ of 0.01, as in my first example, and α of 0.3. Most cases fall in the range 3 % to 4.5 %, as in the first example, but there are 26 cases (rather than 2) below that range and 20 cases (rather than 1) above it. There is more risk than in the first example, but rather less than in the second.

If I now put the value of α up to 0.6, as shown in Figure 5, I get a wider range than when I increased the value of σ. Most cases are in the range 2.5 % to 5 %, but there are 68 out of 10,000 below that range and 60 out of 10,000 above it. To be really sure about being able to meet any particular guaranteed sum assured, I would rather calculate it on the basis of 2 % interest, knowing that there was only a chance of 4 in 10,000 of results falling below that level.

If I put the value of α further up to 0.9, still keeping the value of σ as 0.01, we get a very much wider spread of results, as shown in Figure 6. 156 out of the 10,000 cases, about 1.5 % of them, show a value of $S(20)$ less than 20, implying a negative rate of interest. The worst return falls into the –3.5 % to –3 % band. At the other end, one simulation gives a result better than a uniform 11.5 %.

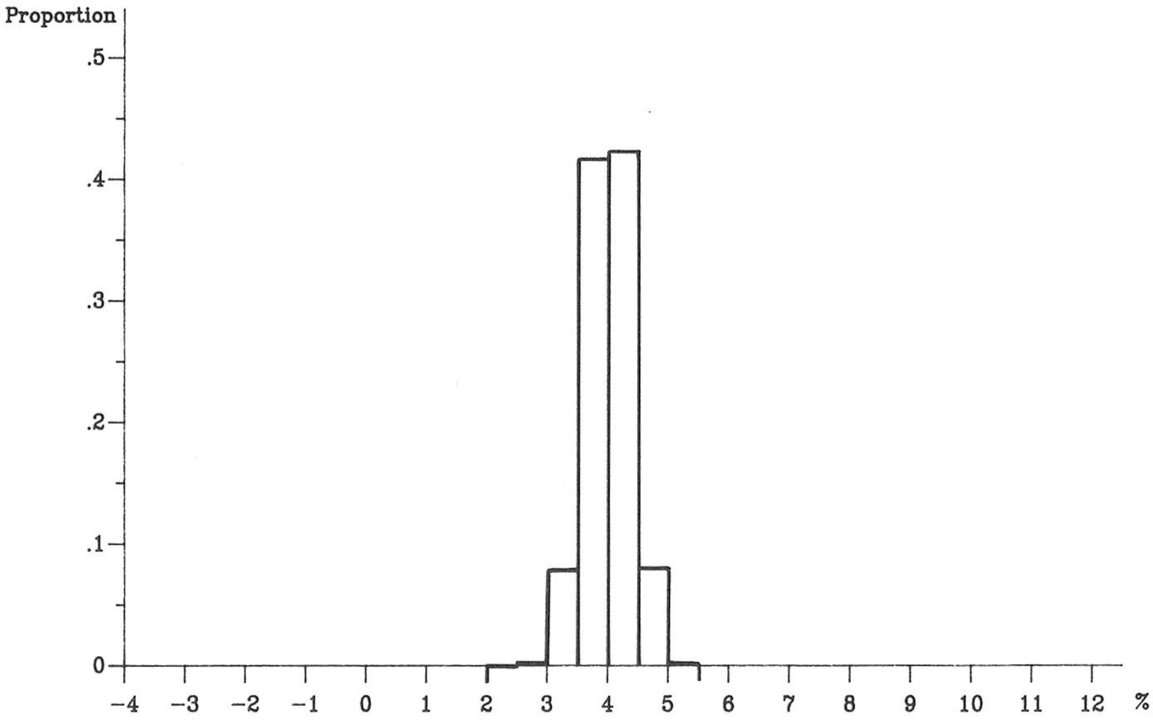

Figure 4: Distribution of $S(20)$; $\alpha = 0.3$, $\sigma = 0.01$

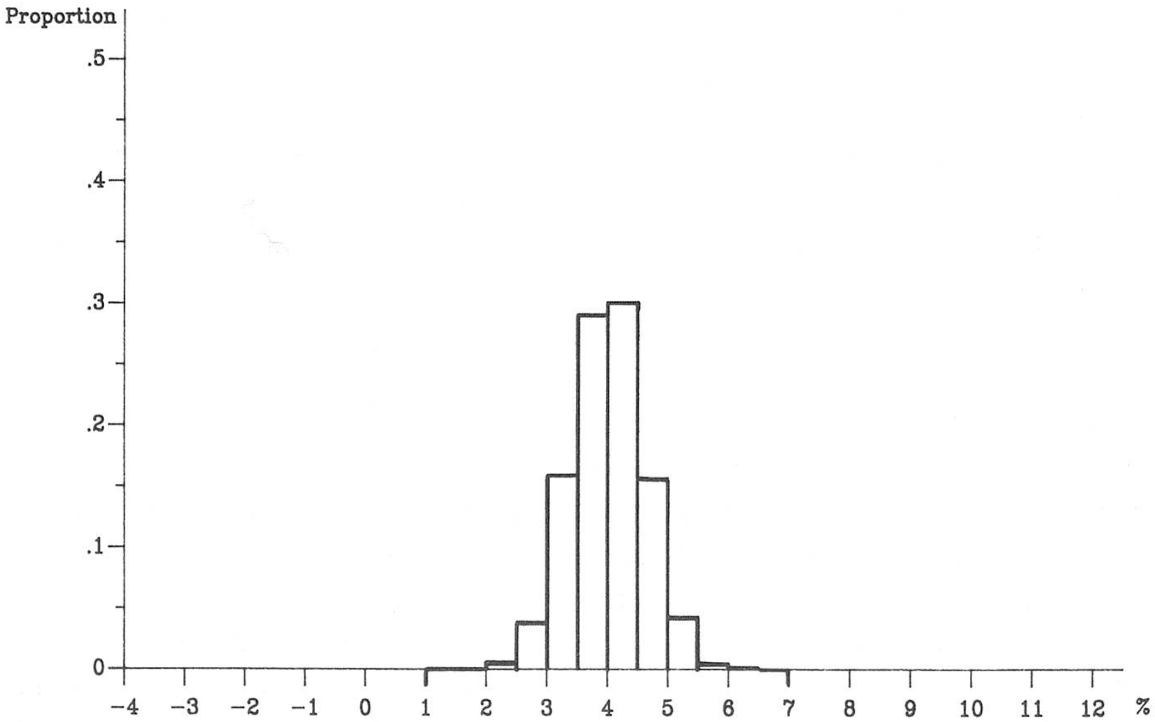

Figure 5: Distribution of $S(20)$; $\alpha = 0.6$, $\sigma = 0.01$

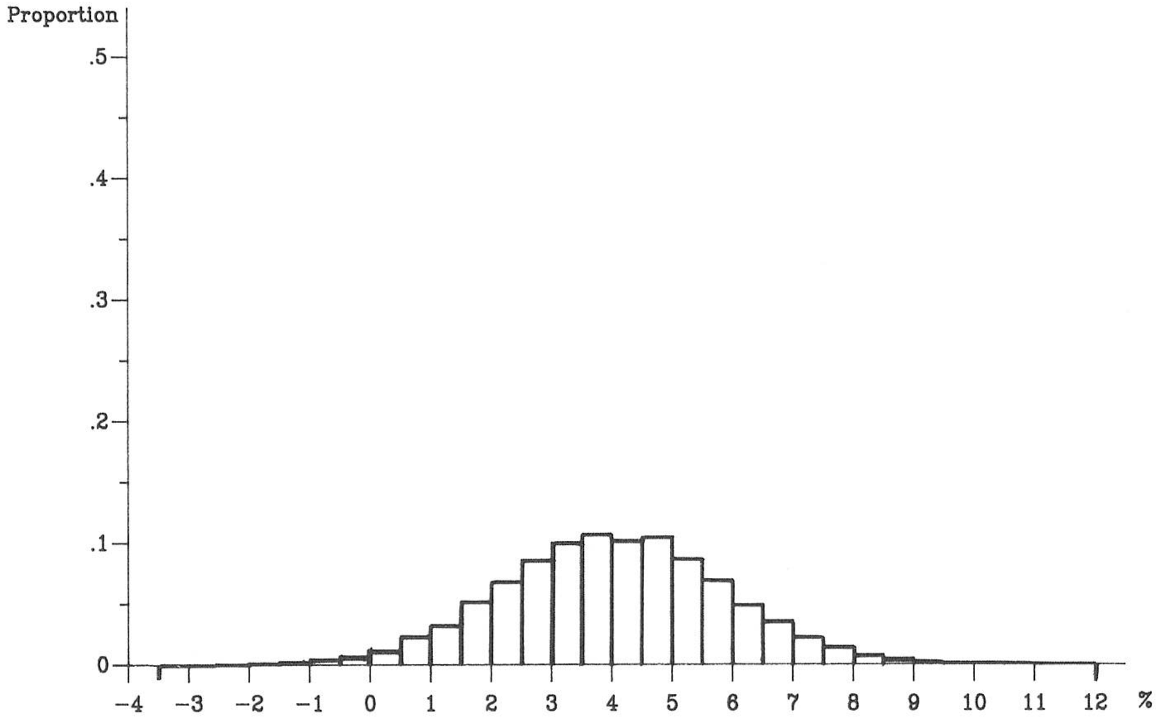

Figure 6: Distribution of $S(20)$; $\alpha = 0.9$, $\sigma = 0.01$

If one is looking for the same sort of security as in the previous examples, one would have to calculate a sum assured on the basis of about –1.5 % interest; even so 16 out of 10,000 simulations produced a result worse than this. But it would not be easy to sell life insurance on a –1.5 % technical interest rate, unless it were sufficiently wrapped up with expenses and mortality to conceal the true interest rate from the policyholder.

Of course, to produce an average negative return there must have been negative rates of interest in several, perhaps many, of the years in the particular simulation. This is possible in my model, because it is not $i(t)$ that is lognormally distributed, implying secure capital and some positive rate of interest; instead, it is $1 + i(t)$ which is lognormally distributed, implying that the total return on an investment of 1 cannot be less than zero, but it may well be less than 1. This reflects the fact that the capital, even of short-term deposits, is not absolutely secure, but depends on the ability of the borrower to repay.

However, I do not want you to draw too many realistic conclusions from this simple example. The point to note is that different assumptions about the stochastic model for investment produce very different results, and in particular that the amount of autocorrelation in the model may be of critical importance.

2 Applications in a life office

I shall now go on to describe in general terms the sort of investigations that one can carry out using some stochastic investment model and some model of a life office, which may be more or less realistic. In general the method used has to be by random simulation. In a few cases it may be possible to get precise analytical results, but I have only found this to be possible with unrealistically simple models.
Any stochastic investment can be used. This could be a simple model such as I have just described, or a more complicated one that includes different categories of asset. Which you use depends on what you want to investigate.
Different model portfolios for the life office can be used. The simplest case is such as I have just described, a single specimen police, perhaps with realistic mortality and expenses included, rather than my pure accumulation.
Another possibility is to use a series of similar policies, issued in successive years, so as to build up a portfolio of policies. You might now wish to allow for growth in new business, and also for lapses and surrenders.
You may wish to assume a mixture of new business in each year, with appropriate proportions of policies of different classes, different terms, and different ages at entry.
At the most extreme you can use the real portfolio of an actual office, possibly simplified by grouping into five-year age groups, rather than by keeping each single age separate. But very often a real life office is too complicated for the sort of investigation that it is worth doing. It seems to me much better to use the simplest model that you need for the particular investigation you wish to carry out. In the example I have described earlier, to put in realistic allowances for mortality and expenses would, I think, just confuse the picture, whereas the practical lesson can be learned, that the particular stochastic investment model makes a difference to the results, with the simplified model I presented.
What are the possible objectives of investigations like this? I suggest that the first may be to determine a premium rate, at least in a market where you are free to do so. If premium rates are fixed by legislation or regulation or by a cartel, then the question may be: is it worth selling policies at these premium rates or not?
In Britain life offices are free to choose their own premium rates, and they are controlled through a system for year-end valuations. Offices are free to offer policies both without profits and with profits. Generally annuities and temporary assurances are without profits, and generally whole of life and endowment assurances are with profits, but there are no rules that say that this must always be the case.

For with profits contracts it is appropriate for an office to calculate premiums on an extremely cautious basis, equivalent to allowing only a small number of simulations out of 10,000 where the premiums are not enough; the excess interest that will be earned in all the other cases can be distributed in the form of bonuses in some way or another.

For without profits contracts it may be reasonable for the office to choose premium rates a bit below the median or the mean return, but not excessively far below. In a majority of cases the office would expect to make a profit, though it also runs a risk of making a loss. The simulation method can then be used to find out how large an additional reserve the office needs in advance, in order to be almost certain that it can pay the guaranteed amounts. It takes extra capital to write without profits business on terms that are reasonably acceptable in the market. Of course, it may be the with profit policyholders that provide this capital, and it may be quite a proper investment for them.

The next sort of investigation one could carry out relates to the methods of bonus distribution. In my simple example, after one year a particular rate of interest would have been earned on the invested assets, and we would have to decide whether or not to declare a bonus, and if so how much to distribute. If we distribute too much, then we run the risk of not being able to pay the increased guaranteed amount, the sum assured plus added bonus; I assume that bonus, as is usual in Britain, is distributed by means of an addition to the sum assured, what the Americans call "paid-up additions". However, if we distribute too little, then our policyholder may be displeased, we may get a poor reputation in the market, and we may not be able to sell enough new business. So the question is: what is the maximum amount of bonus that we can safely distribute?

The traditional method of distributing bonus in Britain was by means of reversionary bonus, increasing the sum assured. Once declared, this was guaranteed, and bonus declarations have moved over the years from quinquennial or triennial to annual. But because of the very large investment by life offices in Britain in ordinary shares (common stocks) whose market values fluctuate considerably, offices have developed a system called "terminal bonus", where a large part of the bonus is distributed at maturity, and is not guaranteed until then. If the life assured dies before the policy reaches maturity, then an appropriate terminal bonus is payable on death. However, if the policyholder chooses to surrender his policy early, often no allowance is made for the terminal bonus, so the penalties for surrender in the later years of a policy may be very great.

In the combined reversionary bonus and terminal bonus system, it is possible to use a simulation method to help to decide how much bonus should be declared in a guaranteed form, and how much should wait until the end, not guaranteed. It may now be sensible to include a realistic allowance for mortality and claims on death, and possibly a realistic allowance for lapses and surrenders.

The next area in which stochastic investigations may be helpful is in determining solvency. In many countries the technical reserves are calculated on the same basis as the premium, but in other countries they are not. In Britain it is the responsibility of the Appointed Actuary of each company to decide on an appropriate valuation basis, which must not show lower reserves than a minimum statutory basis, but which usually is substantially stronger than this minimum basis. A question that insurance supervisors may like to investigate is: what is a suitable statutory basis in order to identify companies which are likely to be inadequate in future years? Or, using a particular method for determining present solvency, what is the probability that a company which just passes the solvency test will prove to be inadequate in some future year?

It may well be that supervisors discover that valuing on the original premium basis, with sufficiently safe premiums, is in fact satisfactory. But it is clear that the present solvency margin prescribed by the First Life Directive of the EEC – in effect 4 % of the technical reserves for all classes of policy, for all ages and all durations – is not equally appropriate to all companies.

So far I have described investigations that could be done assuming only one class of asset, or alternatively that the invested funds of the office behave according to one particular stochastic model. If we allow a number of different classes of asset, there are a great many other investigations that we could carry out. I shall simplify by assuming only two classes of asset: ordinary shares and bonds. I shall assume that when bonds are bought they are of appropriate durations for matching the given liabilities. An alternative is to consider bonds of different durations as different classes of asset, and to investigate the risks involved in not matching correctly by duration.

Let me return to our British position. Many offices in Britain have premiums for with-profit policies that have been calculated a long time ago on the basis of a very low rate of interest, and are therefore very high in relation to the basic sum assured. Such offices are free to invest a high proportion of these premiums, perhaps all of them, in ordinary shares. The question to be asked is: how safe is this investment strategy, bearing in mind the bonus system, a mixture of reversionary and terminal bonus. What is, in some sense, an optimal investment policy, bearing in mind the desire to pay the maximum total bonus to policyholders, subject to the constraint of not becoming insolvent.

It is clear that for a mutual life office maximum bonus to policyholders is an appropriate objective, though quite how that bonus is spread across different generations of policyholder is a further question. For a proprietary office, ie one with shareholders, the objective of course is to maximise the profits to shareholders, but this may best be done by maximising the profits to with-profits policyholders, since shareholders typically receive as profit a fixed proportion of the value of the bonuses to policyholders.

One possible investment strategy would be to invest all the premium in bonds until the guaranteed sum assured has been "bought", ie the investment of all premiums to date is sufficient to guarantee the payment of the sum assured; then future premiums can be invested in as risky a way as one likes. An alternative is to invest all the earlier premiums in ordinary shares, and then, if their investment performance turns out to be poor, to invest the balance of premiums in bonds to meet the guaranteed liabilities. Yet another possibility is to invest in a constant mix of bonds and shares throughout the duration of the policy.

Stochastic simulation allows one to investigate the result of any of these policies.

More realistic, and more complicated, exercises can be constructed, taking into account all the various decisions that could be made by the life office in future, and simulating these appropriately. In particular, the office may be free to choose a particular bonus policy, where the bonus rate may be made dependent on the outcome of each particular stochastic simulation. It is also free to choose an investment policy, where the proportions invested in particular assets may also depend on the outcome from year to year of the stochastic simulation. In each case one has to devise an algorithm capable of being programmed which might represent the decisions made by the office about its bonus policy or about its investment policy consequent on any particular outcome of the simulation so far. Statutory constraints, like remaining solvent, are best included as constraints on what the office can do.

Further elaborations that are possible include making the surrender rates or lapse rates of policyholders also depend on the particular stochastic simulation, or making the new business sold in future years depend on the bonus rates declared by the office. These are complications that I have not seen attempted yet in practice.

However, many of the other things that I have described have been done by actuaries in Britain, mostly in fact in Scotland, and there are papers in *Journal of the Institute of Actuaries* and *Transactions of the Faculty of Actuaries* describing these. Quite a lot of work, however, has been done internally in offices or by consultancies, without being published.

3 The Wilkie stochastic investment model

I now turn to a brief description of one particular stochastic investment model, the one that has become known in Britain as the "Wilkie stochastic investment model". I shall present the model as it appeared most recently in my paper for the Montreal Congress: "Stochastic Investment Models for XXIst Century Actuaries" (*Wilkie*, 1992). In this paper I extended the model from the UK to two other countries, the United States and France.
The data series modelled for each country are: the retail prices index, $Q(t)$; the dividend yield on ordinary shares (common stocks), $Y(t)$; an index of dividends on ordinary shares, $D(t)$; and the yield on long-term or irredeemable bonds, $C(t)$. In each case annual observations of the series have been taken, for the following periods: United Kingdom 1923 – 1990; United States 1926 – 1989; France 1951 – 1989.
For each series I shall state the model, and give rounded "practical" values of the parameters.

Retail prices index

The model for an index of retail prices, $Q(t)$ is straightforward; the force of inflation in each year follows a first order autoregressive process. It can be expressed in a number of stages.
The force of inflation in year $(t-1, t)$ is denoted $I(t)$ and is calculated by

$$I(t) = \ln Q(t) - \ln Q(t-1),$$

which is adjusted to give a value with zero mean

$$IN(t) = I(t) - QMU,$$

which in turn follows an autoregressive process

$$IN(t) = QA \cdot IN(t-1) + QE(t),$$

where the random residual, $QE(t)$, is expressed as

$$QE(t) = QSD \cdot QZ(t),$$

where $QZ(t)$ has zero mean, unit standard deviation, and is assumed to be normally distributed and independent of $QZ(t-k)$, for $k = 1, 2, \ldots$.

Estimated parameter values for the three countries are

	UK	USA	France
QMU	0.05	0.03	0.06
QA	0.6	0.65	0.55
QSD	0.04	0.035	0.04

The similarity in the estimated parameter values is remarkable.

Dividend yield on ordinary shares

The model for dividend yields on ordinary shares, $Y(t)$, is basically a first order autoregressive model for the logarithm of the yield, but an additional influence from inflation is included, on the grounds that high current inflation has an adverse effect on share prices, thus increasing dividend yields.
The model for $Y(t)$ is given by

$$\ln Y(t) = YW \cdot I(t) + YN(t) + YMU,$$

where $YN(t)$ has zero mean, and is given by

$$YN(t) = YA \cdot YN(t-1) + YE(t),$$

and $YE(t)$ is expressed as

$$YE(t) = YSD \cdot YZ(t),$$

where $YZ(t)$ is a sequence of independent unit normal variables.

Estimated parameter values for the three countries are:

	UK	USA	France
YW	1.95	0.5	1.64
YA	0.5	0.7	0.88
$YMU\,\%$	3.8 %	4.3 %	2.4 %
YSD	0.16	0.21	0.165

The parameter values are less close together than are those for retail prices. The low mean yield on ordinary shares in France is not comparable with the yields in the United Kingdom and the United States because it is net of tax, and it should be grossed up at an appropriate rate of tax.

Dividend on ordinary shares

The dividend index on ordinary shares, $D(t)$, is modelled by making changes in dividends depend on the changes in the retail prices index, with a lag, but with unit gain – that is, a 1 % rise in retail prices in due course results in a 1 % rise in dividends – together with some additional terms. The model is

$$\begin{aligned}\ln D(t) = \ln D(t-1) &+ DW \cdot DM(t) + (1-DW) \cdot I(t) + DMU \\ &+ DY \cdot YE(t-1) + DB \cdot DE(t-1) + DE(t)\,.\end{aligned}$$

The term $DM(t)$ represents an exponentially weighted moving average of current and past inflation, and is calculated by

$$DM(t) = DD \cdot I(t) + (1-DD)DM(t-1)$$

and the residual $DE(t)$ is given by

$$DE(t) = DSD \cdot DZ(t)\,,$$

where $DZ(t)$ is a sequence of independent unit normal variables.
The DMU term represent the mean rate of real growth of dividends. The term $DY \cdot YE(t-1)$ represent the observation that changes in share prices anticipate changes in dividends, so share yields fall prior to a dividend rise. The term $DB \cdot DE(t-1)$ represents a carried forward effect of changes in dividends from one year to the next.

Estimated parameter values for the three countries are:

	UK	USA	France
DW	0.8	1.0	1.0
DD	0.2	0.38	0.2
DMU	0.0135	0.0155	0.0
DY	−0.175	−0.35	0.0
DB	0.55	0.5	0.7
DSD	0.06	0.09	0.085

The parameter values have considerably similarity. The response to inflation is similar in all three countries. The mean real rate of growth of dividends in the United Kingdom has been remarkably high in recent years, giving a high value for DMU. In France share prices do not seem to anticipate dividend changes as well as in the other countries.

Yield on long-term bonds

The yield on long-term bonds or irredeemables, $C(t)$, is represented by the "Fisher decomposition" as the sum of an allowance for expected future inflation and a real rate of interest. The allowance for expected future inflation is represented by an exponentially weighted moving average of past inflation. The real rate is represented by a first order autoregressive model plus some additional terms. I had earlier used a third order autoregressive model for the United Kingdom, but this is not necessary for the other two countries, and a first order model seems to represent the UK experience adequately.
The model is thus

$$C(t) = CM(t) + CMU \cdot \exp CN(t)\,,$$

where the moving average of past inflation, $CM(t)$, is calculated by

$$CM(t) = CD \cdot I(t) + (1 - CD)CM(t-1)\,,$$

CMU is the mean real rate of interest, the zero-mean autoregressive part, $CN(t)$, is given by

$$CN(t) = CA \cdot CN(t-1) + CY \cdot YE(t) + CE(t)\,,$$

$CE(t)$ is given by

$$CE(t) = CSD \cdot CZ(t)\,,$$

and $CZ(t)$ is a sequence of independent unit normal variables.
The $CY \cdot YE(t)$ term represents simultaneous movements in share yields and bond yields.

Estimated parameter values for the three countries are:

	UK	USA	France
CD	0.045	0.058	0.2
CA	0.90	0.96	0.90
$CMU\,\%$	3.1 %	2.65 %	2.5 %
CY	0.15	0.07	1.0
CSD	0.175	0.21	0.3

The parameter values are reasonably similar, though in France the "memory" for inflation seems to die away more quickly than in the other countries, the simultaneous influence of share yields is high, and the residual standard deviation is also high.

Comparison with random-walk models

The models described above consider only annual observations. Many investigations have considered share prices or bond yields at much more frequent intervals, daily, weekly or monthly, and have concluded that a pure random walk model fits satisfactorily. This is consistent, in the short term, with my model.
Consider ordinary shares. In the short term dividends on a share index change very slowly, and for an individual ordinary share change only in quarterly or half-year steps. The main short-term determinant of price is the yield. The term in the yield model representing the simultaneous influence of inflation is quite small and changes only very slowly. This leaves a simple first order autoregressive model, with a parameter, taking the US case, of 0.7 over one year. This means that yields revert to their mean level by only 30 % of their deviation from that mean within one year. The model is consistent with a first-order autoregressive model over any shorter timescale, with a different autoregressive parameter, and a different

standard deviation. For a monthly model the autoregressive parameter becomes $0.7^{1/12} = 0.97$. For a weekly model the parameter is 0.993, and for a daily model 0.999. It is not surprising that other investigators have taken these values as equal to unity.

It is clear that in the long run share prices must be dependent on company earnings or company dividends, which themselves must be related. The pure random-walk model leads, in the long run, to no connection at all between share prices and share dividends, which is surely untenable.

A similar argument applies to bond yields. The autoregressive parameter for this yearly model ranges from 0.90 in the United Kingdom and France to 0.96 in the United States. The corresponding parameters for monthly observations are thus very close to unity, and for daily observations even closer. Yet it is unreasonable to suppose that in the long run interest rates move wholly without constraint. The several thousand years of interest rate history recorded by *Homer* (1963) testify to the fact that an interest rate of 3 % is low, and, except in an inflationary environment, an interest rate of 10 % is high.

Integrating the models

What I have described so far gives three independent models for the three countries investigated. For simulation purposes one may want to use them together, combined in an appropriate way. In my recent Montreal paper I described fully how they can be combined, and I do not have time to go into detail today. I shall just outline the method.

I first investigated the lagged and simultaneous cross-correlations between the residuals of each of the four series in the three different countries. Within country lagged correlations are already taken account of in this model. It fortunately turns out that there are no important cross correlations across countries except for simultaneous correlations, and those only for matching series, that is, the residuals of the inflation series in the three countries, the QZs, are simultaneously correlated. These correlations can be dealt with by making the residuals depend on three independent random variates, perhaps in a cascade fashion by calculating the Choleski decomposition of the covariance matrix.

The series for retail prices and for dividends are expressed in local currency – pounds, dollars or francs as appropriate. So in order to do calculations in one currency we need to take the exchange rates between the three currencies into account. Purchasing power parity – the idea that exchange rates reflect changes in retail prices in different countries – forms the basis of my model. But currencies clearly do not fall in line exactly with purchasing power parity, i.e. changes in the respective retail prices indices, and it is convenient first to model the discrepancy between the purchasing power parity index and the actual index by an autoregressive model.

But in addition, arbitrage in the foreign exchange markets means that the cross rates must always balance. With three currencies there are really only two independent exchange rates. This suggested to me that another way of modelling exchange rates would be to postulate a hidden series for each country, a series related to the retail prices index in that country with a suitable delay, and to make the actual exchange rates depend on the ratios of these hidden series. The model is now saturated with unknown parameters, since the hidden series cannot be observed, but with some reasonable assumptions it is possible to construct models for these hidden series which are satisfactory for simulation.

Comparison with other models

I would claim that my investment model is the most comprehensive model published so far that takes into account the long-term stability of share dividend yields and interest rates. The work of *Ibbotson/Sinquefield* (1989) in the United States covers a wider range of investments, but because they mainly work in terms of cumulative wealth ratios they do not in general recognise these long-term stabilities. Their model has therefore substantial short-term uses, but seems less satisfactory for long-term applications.

Desirable extensions

My model remains deficient in a number of respects. Although retail prices are modelled explicitly, salaries and wages are not. It is possible to specify that salaries increase in line with prices, with perhaps a constant differential. But this is less satisfactory than a stochastic model for the differential. Salaries are of importance to pension funds, and possibly have a greater influence on life offices expenses and on the amounts of some general insurance claims than do retail prices.

Another gap in my model is in the representation of the term structure of interest rates. The model includes only one interest rate, the yield on irredeemables. It is possible to assume that short-term rates depend on long-term rates with a constant differential, and define a fixed yield curve between the two, but it would be more satisfactory to represent the whole yield curve on a stochastic basis, as *Tilley* (1990) and *Tilley/Mueller* (1991) have done. Their model, however, does not take into account the influence of expected inflation on interest rates, as does mine. Research is needed here to integrate the two approaches.

My model so far does not consider property (real estate) as a possible investment, though *Daykin/Hey* (1990) postulate an arbitrary model for property on the same lines as my model for shares. Investigation based on real data would be desirable.

The other considerable gap in my original model was in international comparisons. What I have described today begins to cover this gap. An investor in some other country requires a basic model to represent his home market, and then requires investments in other countries modelled in relation to his own.

To cover all important markets, and to extend the series considered for all markets to include the extensions discussed above, would be a considerable task, requiring the accumulation of large amounts of data.

Conclusion

My stochastic investment model forms the basis for a comprehensive long-term stochastic investment model for use by actuaries in modelling future uncertainty. I hope that you in Switzerland will be able to make use of these ideas, and to carry forward yourselves research into this sort of subject. Indeed I might venture to suggest that, in your series of parallel meetings that take place on the day before this Annual Meeting, you might think of including an investment topic, a Swiss AFIR-group, to correspond with your Swiss ASTIN-group.

A. D. Wilkie
R. Watson & Sons, Consulting Actuaries
Watson House, London Road
Reigate, Surrey RH2 9PQ

References

Daykin, C.D./Hey, G.B. (1990): Managing uncertainty in a general insurance company. Journal of the Institute of Actuaries *117*, 173.

Homer, Sidney (1963): A History of Interest Rates: 2000 BC to the present. Rutgers University Press, New Brunswick, New Jersey.

Ibbotson, R.G./Sinquefield R.A. (1989): Stocks, bonds, bills and inflation: historical returns (1926–1987). The Research Foundation of the Institute of Chartered Financial Analysts, Charlottesville, Virginia.

Tilley, James A. (1990): A stochastic yield curve model for asset/liability simulation. Proceedings of the 1st AFIR International Colloquium, Paris 1990, 2, 65.

Tilley, James A./Mueller, Mark (1991): Managing interest rate risk for long liabilities. Proceedings of the 2nd AFIR International Colloquium, Brighton 1991, 1, 169.

Wilkie, A.D. (1986a): A stochastic investment model for actuarial use. Transactions of the Faculty of Actuaries *39*, 341.

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

Wilkie, A.D. (1987): Stochastic investment models – theory and applications. Insurance: Mathematics and Economics *6*, 65.

Wilkie, A.D. (1992): Stochastic investment models for XXIst century actuaries. Transactions of the 24th International Congress of Actuaries, Montreal *5*, 119.

Summary

Several stochastic investment models are presented and their actuarial applications are discussed.

Zusammenfassung

Verschiedene stochastische Investitionsmodelle werden präsentiert und ihre Anwendungen in der Versicherungsmathematik erläutert.

Résumé

Plusieurs modèles stochastiques ont été présentés et leurs applications actuarielles discutées.