JSS 29 (1986) 25-51

SOME APPLICATIONS OF STOCHASTIC INVESTMENT MODELS

BY A. D. WILKIE, M.A., F.F.A., F.I.A.

(A paper presented to the Institute of Actuaries Students' Society on 29 January 1985)

1. INTRODUCTION

IN November 1983 I gave a talk to the General Applications Section of the Royal Statistical Society jointly with the Institute of Actuaries Students' Society, entitled "Time-Series Analysis in Investment Models: Inflation and Share Prices", in which I outlined the statistical derivation of a stochastic investment model. This model was presented to the actuarial profession in a paper "A Stochastic Investment Model for Actuarial Use" discussed at the Faculty of Actuaries in November 1984. The statistical background, and other information, is discussed in Occasional Actuarial Research Discussion Paper No. 36, "Steps Towards a Stochastic Investment Model", which is a sort of Appendix to the Faculty paper. It contains most of the material presented to the Students' Society in 1983.

1.2 The Faculty paper described a number of possible applications of stochastic investment models. I stated there that a "stochastic model for investments . . . can be used by actuaries in almost any circumstance where a rate of interest enters their calculations at present". That is a pretty wide field. But I suggested that it opened up "wider possibilities for investigation too". I intend in this paper to describe a number of simple applications of my stochastic investment model, in order to show the flavour of how such a model can be used. None of the topics has been comprehensively investigated, and I am sure that the reader will soon ask questions which I have not answered. I hope that he will be able to think of ways of answering them himself.

1.3. Since the papers referred to above have not yet been widely published, I state the form of the stochastic model in Part 2 of this paper, without going into details of the justification of it. In Part 3 I discuss a simple application of it to methods of charging for expenses. In Part 4 I discuss a possible dynamic investment policy. In Part 5 I introduce the concept of an index-linked annuity with a money guarantee, and discuss how one might calculate the premium for such a contract; in so doing, I introduce an aspect of profit testing. In Part 6 I describe an alternative way of implementing the model to allow for uncertainty in the parameters. And in Part 7 I give some simple hints about the methods of implementing the programming of such models.

2. THE STOCHASTIC MODEL

2.1 The model involves four separate variables: the Retail Prices Index, $Q(t)$; the

Dividends on ordinary shares, $D(t)$; the Dividend Yield on ordinary shares, $Y(t)$; and the yield on irredeemable fixed-interest stocks, 'Consols', $C(t)$. However, I did not find it necessary to use a full multivariate structure, in which each variable could affect each of the others. Instead I chose to use a 'cascade' one, which can be shown diagrammatically below, where the arrows indicate the direction of influence.

$$
\begin{tikzcd}
 & & \text{RETAIL PRICES INDEX} \arrow[d] \arrow[ddll] \arrow[ddrr] & & \\
 & & \text{SHARE YIELD} \arrow[dll] \arrow[drr] & & \\
\text{SHARE DIVIDENDS} & & & & \text{CONSOLS YIELD}
\end{tikzcd}
$$

2.2 Thus the Retail Prices Index series, $Q(t)$, is described first, entirely in terms of its own previous values, and the values of a random 'white noise' series. White noise is the name given by electrical engineers to a sequence of independent identically distributed random variables, which thus have no single dominant frequency, and so bear the same relation to sound as white light does to light.

2.3 The model for $Q(t)$ is

$$\nabla\ lnQ1(t) = QMU + QA(\nabla\ lnQ(t-1) - QMU) + \mathrm{QSD} \cdot QZ(t),$$

where the backwards difference operator ∇ is defined by

$$\nabla\ X(t) = X(t) - X(t-1).$$

and $QZ(t)$ is a sequence of independent identically distributed unit normal variates.

2.4 This model says that the annual rate of inflation follows a first order autoregressive process, with a fixed mean QMU, and a parameter QA such that the expected rate of inflation each year is equal to the mean plus QA times last year's deviation from the mean. Appropriate values for the parameters are:

$$QMU = \cdot 05,\ QA = \cdot 6,\ QSD = \cdot 05.$$

There is fairly little uncertainty about the appropriate values for QA and QSD, but considerable uncertainty about the value to use for QMU, where anything between $\cdot 04$ and $\cdot 10$ might be justifiable, depending on the past period of observation one wishes to take into account.

2.5 The share yield, $Y(t)$, depends both on the current level of inflation and on previous values of itself and on a white noise series. The model is:

$$lnY(t) = YW \cdot\ \nabla\ lnQ(t) + YN(t),$$

where

$$YN(t) = ln\ YMU + YA(YN(t-1) - lnYMU) + YE(t),$$

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

and $YZ(t)$ is a sequence of independent identically distributed unit normal variates.

2.6 This model says that the natural logarithm of the yield consists of two parts: the first is directly dependent on the current rate of inflation (a high rate of inflation implying a high share yield and vice versa), and the second, $YN(t)$, follows a first order autoregressive model, similar to that of the rate of inflation itself. Appropriate values for the parameters are:

$$YMU = \cdot 04,\ YA = \cdot 6,\ YW = 1{\cdot}35,\ YSD = \cdot 175.$$

2.7 The index of share dividends, $D(t)$, is made to depend on inflation, both with an exponentially lagged effect and an additional direct effect, and on the residual, $YE(t)$, from the yield model, plus a white noise series, which has both a simultaneous and a lagged effect. The parameters are such that a given percentage increase in the Retail Prices Index ultimately results in the same percentage increase in the dividend index, so the model is said to have unit gain. The model is:

$$\nabla\ lnD(t) = DW\left(\frac{DD}{1-(1-DD)B}\right)\nabla\ lnQ(t) + DX\ \nabla\ lnQ(t)$$
$$+ DMU + DY \cdot YE(t-1) + DE(t) + DB \cdot DE(t-1),$$

where the backwards step operator B is defined by

$$B \cdot X(t) = X(t-1),$$
$$DE(t) = DSD \cdot DZ(t),$$

and $DZ(t)$ is a sequence of independent identically distributed unit normal variates.

2.8 The term in parentheses above involving DD represents an infinite series of lag effects, with exponentially declining coefficients:

$$DD,$$
$$DD\,(1-DD),$$
$$DD\,(1-DD)^2,$$
etc.

The sum of these coefficients is unity, so this part of the formula represents the lagged effect of inflation, with unit gain. This means that if retail prices rise by 1% this term will also, eventually, cause dividends to rise by 1%.

2.9 We can alternatively describe it as the 'carried forward' effect of inflation, $DM(t)$, where

$$DM(t) = DD.\ \nabla\ ln\ Q(t) + (1-DD)\ DM(t-1),$$

from which we see that the amount that enters the dividend model each year is DD times the current inflation rate, plus $(1-DD)$ times the amount brought forward from the previous year, and that this total is then carried forward to the next year.

2.10 Appropriate values for the parameters are:

$$DW = \cdot 8,\quad DD = \cdot 2,\quad DX = \cdot 2,\quad DMU = \cdot 0,\quad DY = -\cdot 2,\quad DB = \cdot 375,\quad DSD = \cdot 075.$$

2.11 The model makes the dividend index appear to depend on the residual of the share yield. In fact share prices to some extent correctly anticipate changes in dividends. For example, an unusual rise in dividends may be correctly forecast by investment analysts, so that share prices take account of this and so rise. The yield is calculated on the previous year's dividend, and so falls. Although this is the *causal* sequence, it is convenient in the model to reflect the *temporal* sequence, so that an unexpected fall in yields results in an upwards change in the dividend index in the following period.

2.12 Although the parameter DMU is set to zero, it is retained in the model, since one may wish to investigate the results of assuming a small positive or negative value for it, implying a positive or negative long-term change in real dividends.

2.13 The Consols yield, $C(t)$, is assumed to consist of a real part, $CN(t)$, plus an allowance for expected future inflation. The latter is based on the actual values of present and past inflation. The real part is defined by a third order autoregressive model, together with an influence from the residual of the yield series, $YE(t)$, and a residual white noise series. The model is:

$$C(t) = CW\left(\frac{CD}{1-(1-CD)B}\right)\nabla\ lnQ(t) + CN(t),$$

where

$$\begin{aligned} lnCN(t) = {} & lnCMU + \\ & (CA1\cdot B + CA2\cdot B^2 + CA3\cdot B^3)(lnCN(t) - lnCMU) \\ & \qquad + CY\cdot YE(t) + CSD\cdot CZ(t), \end{aligned}$$

where $CZ(t)$ is a sequence of independent identically distributed unit normal variates.

2.14 The term in parentheses in CD has a similar form to the DD term in the dividend model, though the parameter value is different. It represents the current value of expected future inflation as an exponentially weighted moving average of past rates of inflation.

2.15 Appropriate values for the parameters are:

$$\begin{gathered} CW = 1\cdot 0,\ CD = \cdot 045,\ CMU = \cdot 035,\ CA1 = 1\cdot 20, \\ CA2 = -\cdot 48,\ CA3 = \cdot 20,\ CY = \cdot 06,\ CSD = \cdot 14. \end{gathered}$$

The value of CW is 1·0, and it might appear that this term could be omitted; however, it may be of interest to investigate variations in this parameter.

2.16 This form of the model says that the influence of inflation on the Consols yield is reflected by using as expected inflation an exponentially weighted moving average of past inflation, with a parameter of ·045. The real rate of return has a mean of 3·5%, and follows a third order autoregressive series with a principal factor of about ·91, so that it tends back towards its mean rather slowly.

2.17 It will be seen that the complete model is wholly self-contained. The only inputs are the four separate white noise series, and no exogenous variables are included. In my view, whatever may be the case for short-term forecasting, such a self-contained model is better for long-term simulations. The rate of inflation, the amount of company dividends, the level of interest rates, and the prices at which shares trade may well depend on such extraneous factors as government policy, business conditions and the political, military, economic and climatic condition of the world. Wars, famines and natural disasters may or may not occur. But they are not forecastable in the long run and their influence is subsumed in the white noise series.

2.18 It would be possible to derive analytically the joint probability distribution of the unknown values of certain of the variables in successive future years, given a suitable set of data to represent the past history and current state at some particular starting time. However, it seems to me particularly complicated to do this for any realistic actuarial purpose, whereas a simulation method facilitates many more possible investigations. The method of simulation that is appropriate for this model is similar to that used by the Maturity Guarantees Working Party. On the basis of a starting position at time $t = 0$, one can generate values for the four series, $Q(t)$, $Y(t)$, $D(t)$, and $C(t)$, for $t = 1$ to N, where N is for example 100. It is necessary to simulate independent unit normal pseudo-random variables for each of the white noise series, QZ, YZ, DZ and CZ, using for example Marsaglia's Polar method, as described in Appendix E of the MGWP Report.

2.19 It is necessary to choose certain initial values to represent the present state, and to start the indices. One can set $Q(0)$ arbitrarily as 1. The model for the Retail Prices Index requires us to postulate a value for $\nabla\ lnQ(0)$, the rate of inflation 'last year', i.e. in the year just preceding the beginning of the simulation period. I denote this by QI. A neutral value for QI is QMU, the average force of inflation. Howeever, one may wish to investigate the effect of a different starting value, or to insert the actual current real value.

2.20 The model for yields requires us to choose a value for the share yield at the start of the simulation period. This is $Y(0)$ or YI. A neutral value for this is given by $YMU \cdot exp\ (YW \cdot QMU)$. As with inflation, it may be of interest to investigate the effect of choosing different values for the starting yield, such as the actual current value. The model for yields requires also a value for $\nabla\ lnQ(0)$, which has already been given by QI.

2.21 To start the dividend series one needs to choose an arbitrary value for $D(0)$. It is of no importance whether one uses a value of 1, or a value equal to $Y(0)$, which would then imply a starting share price, $P(0)$, of 1; either may be used. One then needs to choose a value for the carried forward exponentially lagged effect of inflation, viz:

$$DM(0) = \left(\frac{DD}{1-(1-DD)B}\right) \nabla\ lnQ(0),$$

which I denote just as DM. The neutral value for this is also QMU, but one may wish to use an estimate of the current carry forward. One also needs a value of $\nabla\ lnQ(0)$, given as before by QI. One then needs a value for $YEI = YE(0)$, the random residual that took the share yield to its present level. This could either be stated explicitely, or be calculated given also values for $Y(-1)$ and $\nabla\ 1nQ(-1)$. The neutral value is zero.

2.22 The starting values required for the Consols yield series include a carry forward from past inflation, similar to that required for dividends, though based on a different parameter, viz:

$$\left(\frac{CD}{1-(1-CD)B}\right)\nabla\ lnQ(0),$$

which I denote CM. The neutral value for this is QMU. One also needs to select values for the starting Consols yield, $C(0)$, and for the two past years, $C(-1)$ and $C(-2)$. The neutral value for these is $QMU + CMU$, but the actual current values could be used. The model for $C(t)$ would allow the possibility of negative values of the yield if inflation were negative for long enough. To avoid these occurring I postulate a minimum value for $C(t)$ of $CMIN$, set equal to ·5%.

2.23 Besides calculating values for the four basic series it is also convenient to calculate values for three derived series. The first of these is the share price, $P(t)$, which is easily derived from the formula:

$$P(t) = D(t)/Y(t).$$

2.24 One can next calculate a 'rolled-up' share index being the value of a share index where dividends, net of tax, are reinvested in shares. This is denoted $PR(t)$, where

$$PR(t) = PR(t-1)\left(\frac{P(t) + D(t)\,(1-tax\ A)}{P(t-1)}\right),$$

and $tax\ A$ is the rate of tax on share dividends, assumed constant. In fact I have taken this normally as zero, so I have assumed a gross roll-up. An arbitrary starting value of $PR(0) = 1$ is appropriate.

2.25 The third additional series is a corresponding rolled-up index for Consols, denoted $CR(t)$, where

$$CR(t) = CR(t-1)\left(\frac{1}{C(t)} + (1-tax\ B)\right)C(t-1),$$

and $tax\ B$ is the rate of tax on 'unfranked' income. I take this also normally as zero. An arbitrary starting value of $CR(0) = 1$ is appropriate. This formula assumes that 'Consols' are truly irredeemable stocks, and would not be repaid if interest rates fell below the coupon rate, possibly being then refinanced at a lower coupon rate. This complication could easily be allowed for in the calculation of $CR(t)$ if desired.

2.26 The model described above is called in the Faculty paper the Full Standard Basis, and I have used it throughout what follows. Table 2.1 shows the values of the parameters for this basis, and Table 2.2 shows the results of 1,000 simulations each for 100 years. These are the same as appeared in the Faculty paper. In that paper I also described a Reduced Standard Basis, which gave similar results to the Full Standard Basis, with a slightly smaller number of

Table 2.1. *Values of Parameters in Standard Bases*

	Full Standard Basis	*Reduced Standard Basis*
Inflation:		
QMU	·05	·05
QA	·6	·6
QSD	·05	·05
Share Yield:		
YW	1·35	1·35
YMU%	4·0	4·0
YA	·6	·6
YSD	·175	·175
Share Dividend:		
DW	·8	·8
DD	·2	·2
DX	·2	·2
DY	−·2	−·3
DMU	·0	·0
DB	·375	·0
DSD	·075	·1
Consols Yield:		
CW	1·0	1·0
CD	·045	·05
CMU%	3·5	3·5
CY	·06	·0
*CA*1	1·20	·91
*CA*2	−·48	·0
*CA*3	·20	·0
CSD	·14	·165
Initial Conditions:		
QI	·05	·05
YI%	4·27932	4·27932
YEI	·0	·0
DEI	·0	·0
DM	·05	·05
CI%	8·5	8·5
CM	·05	·05

Table 2.2. *Results on Full Standard Basis*

Term (years)	*1*	*5*	*10*	*15*	*20*	*30*	*50*	*75*	*100*
Mean Rate of Inflation (*GQ*):									
E(GQ)	5·37	5·37	5·14	5·16	5·07	5·03	4·99	5·06	5·08
SD(GQ)	5·34	4·48	3·61	3·13	2·83	2·38	1·80	1·54	1·34
Mean Rate of Money Return on Shares (*GPR*):									
E(GPR)	12·17	10·17	10·07	9·85	9·80	9·69	9·66	9·72	9·71
SD(GPR)	21·72	7·27	5·23	4·28	3·78	3·29	2·60	2·14	1·94
Correlation Coefficient:									
C(GPR,GQ)	−·26	·24	·46	·57	·64	·71	·76	·77	·78
Mean Rate of Money Return on Consols (*GCR*):									
E(GCR)	8·05	8·53	8·69	8·68	8·73	8·74	8·71	8·66	8·68
SD(GCR)	6·27	2·81	1·55	1·12	1·00	1·08	1·19	1·15	1·13
Correlation Coefficients:									
C(GCR,GQ)	−·44	−·62	−·61	−·37	−·13	·27	·58	·67	·73
C(GCR,GPR)	·18	−·09	−·17	−·09	·02	·29	·49	·54	·60
Mean Rate of Real Return on Shares (*JPR*):									
E(JPR)	6·99	4·68	4·73	4·48	4·52	4·44	4·44	4·44	4·41
SD(JPR)	22·75	7·29	4·59	3·46	2·82	2·20	1·61	1·29	1·17
Correlation Coefficient:									
C(JPR,GQ)	−·47	−·39	−·28	−·23	−·18	−·06	·05	·04	·09
Mean Rate of Real Return on Consols (*JCR*):									
E(JCR)	2·94	3·26	3·53	3·45	3·56	3·58	3·56	3·44	3·43
SD(JCR)	9·53	6·41	4·62	3·63	3·07	2·31	1·46	1·13	·91
Correlation Coefficients:									
C(JCR,GQ)	−·82	−·94	−·96	−·96	−·95	−·90	−·77	−·69	−·59
C(JCR,JPR)	·43	·39	·31	·27	·21	·12	·02	·01	·01

parameters. The parameters for it are recorded in Table 2.1, but I have not discussed them above, nor do I use this basis in what follows.

3. UNIT TRUST EXPENSES

3.1 The Managers of Unit Trusts usually make two forms of charge to cover their expenses, a Preliminary Charge, paid when a unitholder purchases Units, and an annual (or periodic) Management Charge, which is usually based on the asset value of the Fund. The first of these charges is collected at the same time as the corresponding expenses, such as commission, are incurred, and I shall not be concerned with it any further. My question is: is a charge expressed as a percentage of the assets likely to be sufficient to meet recurring expenses, which are likely to be closely correlated with the general level of prices? And secondarily: is it better to express the expenses as a percentage of the asset value or as a percentage of the income of the Fund?

3.2 In this exercise I consider only a Unit Trust which is invested in ordinary shares, which have exactly the same investment performance as the ordinary shares in my stochastic model. I ignore the costs of buying and selling shares, and of reinvesting dividend income. I consider both Income Units, where the income, less charges, is fully paid out to the unitholder each year, and Accumulation Units, where the income, less charges, is reinvested in the Fund. Accumulation Units correspond with the usual sort of internal fund of a life office that sells linked policies and reinvests the income. Income Units do not quite correspond with Capital Units in a linked policy, since with these the whole of the income may be made available to the life office to defray its initial expenses.

3.3 I assume that tax is charged on the interest income, less charges, at 30%, the current rate on franked investment income; if the charges were to exceed the income my program would allow an immediate tax credit on the excess of charges. I do not know whether this has in fact happened in the simulations I have carried out. (It wouldn't be difficult to test for this in the program; I just haven't done it.)

3.4 I assume that all expenses are incurred and charges levied at the end of each year; also that all dividends are received at the end of each year. I assume that the Units are actually held for ten years. I do not know whether this is a realistic holding period for actual Units, but it corresponds with a frequently found term for linked policies.

3.5 I assume that the expenses of management start at 1% of the initial amount invested, and change proportionately to the Retail Prices Index, $Q(t)$ of my model. For an initial investment of £1,000 this would mean an annual expense at time 0 of £10, indexed subsequently to prices. In fact this £10 will not be incurred, since the first expenses are assumed to be incurred at the end of the first year, by which time we are looking at the value of $Q(1)$, rather than $Q(0) = 1{\cdot}0$.

3.6 I assume that charges are levied in two alternative ways: (A) 1% of the asset value, and (B) 23·368% of the gross annual income. The charge for Income Units is levied on the capital value of the Units, excluding dividends, which is given simply by the price of shares, $P(t)$ in my model. At time 0 this charge would exactly meet the expenses, though by the end of the first year the price of the shares may well have changed in a different way from the Retail Prices Index. In my simulations I choose a neutral position for the initial yield on shares of 4·27932%, and 23·368% of this is equal to 1% of the initial investment so again at time 0 this charge would exactly meet the expenses. However at the end of the first year the dividends paid are in fact $D(1)$ from my model, and the charge levied is 23·368% of this. Such a charge based on dividends clearly cannot exceed the dividend income; I have (carelessly) made no check that the charge based on assets does not exceed the dividend income; this would be possible if yields fell as low as 1%, which is unlikely, but not impossible from the model.

3.7 The charges for Accumulation Units are the same percentages, but the charge based on assets is calculated on the price including dividend income at the end of each year, and the income net of charges and tax is assumed to be reinvested, so that both the asset value and future income are increased from this cause, though they also fluctuate because of the model. The charges levied on Accumulation Units must be higher than those levied under Income Units, and must be expected to increase.

3.8 For the simulations I have used the Full Standard Basis described in Part 2 of this paper, with the neutral initial conditions. Each of 1,000 simulations starts from the same neutral position. Each simulation is for 10 years, and each produces 10 values for the expenses incurred, and 10 values for each of the two methods of charging for each of Income and Accumulation Units. We must find some way of summarizing these 50,000 numbers. The first way is simply to total the annual expenses and annual charges, and compare the totals; in effect they are discounted at 0% interest. Within a stochastic investment model it is not always clear what meaning can be attached to conventional actuarial discounting at a fixed rate of interest. Yet when comparing two sequences of payments over time there is a natural actuarial desire to compare them in terms of present values. I have therefore calculated the discounted present values of each of the streams of expenses and charges at an interest rate of 6·0%, which corresponds roughly with the net of tax equivalent of the mean yield in my model on fixed interest stock of 8·5%. It should be noted, however, that the differences between any two of my streams of payments may change sign several times over the period of 10 years, so that discounting at different rates of interest may change the sign of the differences in present values in different ways.

3.9 Table 3.1 gives certain statistics for the totals of expenses and charges on the four different bases, and for the discounted present values of these, per £1,000 initial investment. The statistics shown in each case are: mean, standard

Table 3.1

Totals of expenses and charges for 10 years

	Mean	*Standard deviation*	*Skewness*	*Kurtosis*	*Quantiles 2·5%*	*97·5%*
Expenses	£137.8	31·5	·9	3·9	85	215
% Price I.U.	£144.6	44·2	1·2	5·2	75	265
% Div. I.U.	£140·1	39·1	1·1	5·0	85	235
% Price A.U.	£167.2	51·0	1·2	5·3	85	305
% Div. A.U.	£157.7	45·5	1·2	5·2	95	265

Discounted present values of expenses and charges for 10 years

	Mean	*Standard deviation*	*Skewness*	*Kurtosis*	*Quantiles 2·5%*	*97·5%*
Expenses	£ 98.6	20·8	·8	3·8	65	145
% Price I.U.	£103.3	29·6	1·1	5·0	55	185
% Div. I.U.	£100.2	25·9	1·1	4·8	65	165
% Price A.U.	£118.2	33·7	1·1	5·0	65	205
% Div. A.U.	£111.5	29·9	1·1	5·0	65	185

deviation, skewness coefficient ($\sqrt{\beta_1}$), kurtosis coefficient (β_2) and the approximate 2·5% and 97·5% quantiles, which therefore give a symmetric 95% forecast interval.

3.10 As expected, the charges for Accumulation Units are higher than for Income Units. The mean charge based on price for Income Units is higher than the mean charge based on dividends for Income Units, which in turn is higher than the mean expenses. This is true also when the amounts are discounted. But the standard deviations of the higher amounts are larger, and the spread is somewhat greater. The distributions are noticeably skew and fat-tailed. (For a normal distribution the skewness coefficient would be 0, and the kurtosis coefficient 3·0.) This feature results from the logarithmic elements in the model. Crudely, if there is an equal chance of halving and doubling, the mean is 1·25, not 1; thus a high standard deviation in itself produces a high mean value. The median values are much closer together, being respectively £134, £139, £133, £160 and £150 for the totals, £96, £100, £96, £113 and £107 for the discounted values.

3.11 We can see this in a different way by showing the number of occasions out of the 1,000 simulations that expenses (both in total and discounted) exceeded the charges on four different bases. These figures are shown in Table 3.2.

Table 3.2. *Number of occasions out of 1,000 that expenses exceeded charges*

	Total	*Discounted*
% Price I.U.	451	451
% Div. I.U.	473	473
% Price A.U.	270	273
% Div. A.U.	288	298

We see that in nearly half the cases the expenses on Income Units were insufficient to meet the charges, and that this was true for over one-quarter of the cases for Accumulation Units.

3.12 We can also count, for each simulation, the number of years in which expenses exceeded the charges on each of the four different bases. The frequency distribution for these counts is shown in Table 3.3. We see for example that there were 115 cases out of 1,000 in which expenses exceeded a charge of 1% of the price of Income Units is none of the 10 years; in 94 of the cases expenses exceeded the charge in one out of the 10 years; and so on down the column till we find 95 cases in which expenses exceeded the charges in every year out of the ten.

3.13 It is interesting to see that the distribution of these numbers of years when the charge is based on the price of Income Units is fairly uniform, whereas when the charge is based on dividends the distribution becomes U-shaped, with peaks at each end. For Accumulation Units there is a tilt towards the lower values, that is the charges are more likely to exceed the expenses; but there is still an uncomfortable peak at the other end for a charge based on dividends for

Table 3.3. *Frequency of number of years out of 10 in which expenses exceeded charges*

Number of years	*% Price I.U.*	*% Div. I.U.*	*% Price A.U.*	*% Div. A.U.*
0	115	208	204	313
1	94	83	144	125
2	97	65	124	85
3	90	67	80	73
4	83	67	98	53
5	95	54	77	57
6	76	58	72	42
7	92	58	62	51
8	71	74	46	44
9	92	81	56	52
10	95	185	37	105
Total	1,000	1,000	1,000	1,000

Accumulation Units, where there were 105 cases out of 1,000 when the expenses exceeded the charges in each of the ten years.

3.14 We see from this that if the basis for charging is chosen so as to break even in year 0 of the investigation, there is quite a high probability, even for Accumulation Units, that the charges will be insufficient, on the basis of all the assumptions I have made. Further, there is quite a high chance that over ten years we do not break even 'on average', but that the business is unprofitable year after year. Note that the distributions shown in Table 3.3 are not at all binomial, as might have been the case if the chances of making a profit or loss in each year were constant and independent.

3.15 The next step in an investigation of methods of charging might be to calculate the ratio of total charge to total expenses over the ten years in each simulation, and record a frequency table of these ratios. One could then choose a probability level, so that the charge would be sufficient to cover the expenses on say 80% of occasions. This would allow one to calculate a percentage charge to achieve this end. Simulations could then be rerun yet again, using the new percentages, in order to obtain statistics and frequency distributions on this new basis. I have not carried this out.

3.16 In real Unit Trusts nowadays there is usually an upper limit of Management Charge fixed in the Trust Deed, but a lower level is actually charged. The Managers have the right to increase the charge, with notice to the Unitholders, up to the maximum in the Deed. To go further than this would require the approval of Unitholders. The procedure I have described above might give a way of fixing the maximum charge in the Deed, so as to reduce the probability of having to ask Unitholders for an increase. One could then make the actual charges appropriate from time to time to give a reasonable level of profit. Similar considerations may apply to internal funds.

3.17 This investigation at least has shown that it is inappropriate to fix the basis of expense charges in advance, without the possibility of change. To be safe, the expenses would need to be unreasonably (and probably uncompetitively) high. If they were pitched at a reasonable level, there would be too high a chance of their being insufficient.

3.18 I had started this investigation thinking that a charge based on dividends would be more stable than one based on price. This is to some extent true, though to a much smaller extent than I had expected. Indeed, a charge based on dividends appears to have a higher chance of being wholly disastrous than one based on price. I make no comment on the marketing implications of a charge expressed as roughly one-quarter of the income as opposed to 1% of the asset value, except to remark that they are roughly the same amount!

4. A DYNAMIC INVESTMENT POLICY

4.1 One of the main lines of argument of what I call 'classical' financial economics is that markets are efficient, in the sense that the prices of securities are always such that no exceptional profits can be made simply by a study of the history of those prices. My own stochastic model, in which there are rather complicated correlations between the returns on shares and fixed interest stock and their previous returns, suggests that the expected returns on different securities are not always the same, or even relatively the same. The expected returns on shares and on 'Consols' over some future number of years depend in part on the present level of yields on each and on their recent history. It is therefore worth investigating whether a 'dynamic' investment policy might give better returns than a static 'buy and hold' one. By 'dynamic' I mean a policy that takes into account the actual progress of prices in the course of each simulation.

4.2 If one is going to assume a policy that involves 'switching' between shares and Consols possibly each year, it is desirable to take into account the transaction costs of switching. I therefore take the opportunity of investigating the effect of transaction costs on the obvious 'buy and hold' strategies too. And, since interest on Consols and dividends on shares are always paid in cash, I allow for the possibility of 'cash' as a security too. I arbitrarily assume that the rate of interest on cash is three-quarters of the yield on Consols, and that the expected future yield on cash is always the same as this. I have no evidence to support this figure of three-quarters, which may well be really too low. However, I have not been able to investigate the behaviour of short-term interest rates relative to long-term ones, and for this exercise I simply wanted a cautious, but not wholly ridiculous value.

4.3 I allow for expenses, both actual and expected, on the following basis:

£100 of cash buys £97 of shares at middle market prices;
£100 of shares at middle market prices buys £98 of cash;
£100 of cash buys £99 of Consols at middle market prices;
£100 of Consols at middle market prices buys £99 of cash.

4.4 I start the exercise with £100 of cash, and at the end of twenty years I convert the current investments back into cash. As already noted each year's income is received in cash and converted into shares or Consols subject to the expenses mentioned above.

4.5 In this example I ignore taxation.

4.6 In Table 2.2 above I quoted the statistics of the rate of return on shares and on Consols over various periods including a 20-year one. These rates of return did not allow for expenses, but assumed that dividends and interest income could be reinvested in shares or Consols free of cost. The first results from the present exercise compares the 20-year returns allowing for expenses, which I shall call 'net' returns, with these earlier 'gross' ones. (Note that I mean net of expenses, not net of tax.) Since holding cash is assumed to involve no expenses, gross and net returns on cash are the same.

4.7 Table 4.1 shows some of the statistics of the distribution of the proceeds of £100 invested at the beginning of 20 years in the securities shown, which are then converted into cash at the end of that period. The lower part of the table shows statistics of the mean compound annual rates of return corresponding to those proceeds. (Note that the gross returns for shares and Consols differ slightly from those shown in Table 2.2 above, because these simulations are taken over a different period, in effect 1,000 consecutive periods of 20 years, whereas the former statistics were measured over the first 20 years of 1,000 consecutive simulations of 100 years each. Different pseudo-random numbers therefore entered the process.) The statistics shown in each case are again: mean, standard

Table 4.1

Proceeds at the end of 20 years from initial investment of £100

	Mean	*Standard deviation*	*Skewness*	*Kurtosis*	*Quantiles 2·5%*	*97·5%*
Shares (gross)	£818	613	2·0	8·8	165	2485
Shares (net)	£760	569	2·0	8·7	155	2305
Consols (gross)	£538	107	1·2	6·0	375	775
Consols (net)	£520	103	1·2	6·0	355	745
Cash	£357	66	·8	4·5	245	495

Mean compound annual rate of return per cent of above proceeds

	Mean	*Standard deviation*	*Skewness*	*Kurtosis*	*Quantiles 2·5%*	*97·5%*
Shares (gross)	9·82%	3·85	·0	2·8	2·6	17·4
Shares (net)	9·41%	3·83	·0	2·8	2·3	17·0
Consols (gross)	8·68%	1·03	·4	3·7	6·8	10·8
Consols (net)	8·50%	1·02	·4	3·7	6·7	10·6
Cash	6·49%	·97	·2	3·3	4·7	8·3

(*Note:* net means net of expenses; taxation is ignored.)

deviation, skewness coefficient ($\sqrt{\beta_1}$), kurtosis coefficient (β_2) and the approximate 2·5% and 97·5% quantiles, which therefore give a symmetric 95% forecast interval.

4.8 It can be seen that the mean net proceeds on shares of £760 are some £58 less than the mean gross proceeds, corresponding to a mean percentage rate of return of 9·41% as compared with the gross 9·82%. For Consols the mean net proceeds are £520, some £18 less than the mean gross proceeds of £538, corresponding to mean rates of return of 8·50% and 8·68% respectively. Cash gives a distinctly poorer mean return, because of the artificially low way I have defined the yield on it.

4.9 It should be noted that the 95% range for share proceeds is very much wider than that for Consols. Further, although the mean rates of return are approximately normally distributed (for a normal distribution the skewness coefficient would be 0 and the kurtosis 3), the distribution of proceeds is very markedly skew and fat-tailed. The median proceeds are very much less than the mean (£605 and £655 for shares net and gross respectively; £515 and £525 for Consols net and gross respectively).

4.10 The net proceeds from shares exceeded that from Consols in 613 out of the 1,000 simulations, and in the remaining 387 the proceeds from Consols exceeded those for shares. The proceeds from cash never came top of the three. (I didn't count whether cash always came bottom; over a short period it might quite well beat either shares or Consols because of changes in the market value of the securities.)

4.11 We now have a standard with which to compare the proceeds from any dynamic investment policy. The one I chose is simple: at the beginning of each year of the simulation I have got certain amounts already invested in cash, shares and Consols. At the beginning of the first year of each simulation I have £100 in cash, zero in the other two. I 'look ahead' a certain number of years, say t years, taking into account the current conditions as the starting point. I then calculate, not the expected returns, but the median proceeds for each of the three possible investments over the t years. I calculate the median proceeds by assuming that all the future random elements, QE, YE, DE and CE, are zero. This calculation does not give the mean values of the proceeds because of the skewness. (I am not even quite sure that it gives the median values, because of the reinvestment of income; but at least it is easy to do!)

4.12 I then consider for each of the possible switches, including a 'hold', what the proceeds would be at the end of t years, allowing for the expenses first of switching (nil for a hold) and then of converting into cash at the end of t years. This is somewhat arbitrary, since we don't have to switch into cash at the end of each t years, but only at the end of the 20-year horizon. Also, I have made no allowance for expenses in the assumed reinvestment of dividends and interest. But at this stage precision is not essential. I am only looking for a criterion on which to base investment policy; the actual outcome of the investment policy is then calculated accurately.

4.13 I said that I look ahead a certain number of years. I have chosen this period in two ways: first, t is a fixed 1, 2, 3, 4 or 5; secondly, t is the number of years remaining to the end of the 20-year horizon. The results are shown in Table 4.2. We soon see that to look ahead for 3 years gives the highest mean proceeds and mean rate of return, though to look ahead for 2 years gives a slightly lower standard deviation with a narrower 95% range. Looking ahead to the end of the 20-year horizon is somewhat less good. The mean proceeds in all cases are some 50% higher than the mean proceeds on shares, net of expenses, shown in Table 4.1. The mean annual rate of return is some 3% higher than on shares. On the face of it, this simple investment policy looks as if it should pay off handsomely.

Table 4.2

Proceeds at the end of 20 years from initial investment of £100, following dynamic investment policy

Looking ahead	*Mean*	*Standard deviation*	*Skewness*	*Kurtosis*	*Quantiles 2·5%*	*97·5%*
1 Year	£1203	767	2·6	15·3	365	3225
2 Years	£1207	716	2·1	9·5	385	3155
3 Years	£1226	755	2·2	10·1	375	3395
4 Years	£1210	752	2·2	10·3	355	3255
5 Years	£1203	750	2·2	10·5	335	3185
to end of 20 years	£1160	750	2·3	11·5	315	3165

Mean compound annual rate of return per cent of above proceeds

Looking ahead	*Mean*	*Standard deviation*	*Skewness*	*Kurtosis*	*Quantiles 2·5%*	*97·5%*
1 Year	12·40%	3·10	·3	3·7	6·6	18·9
2 Years	12·50%	2·96	·3	3·1	6·9	18·8
3 Years	12·53%	3·06	·2	3·2	6·8	19·2
4 Years	12·44%	3·11	·2	3·2	6·4	19·0
5 Years	12·40%	3·12	·2	3·2	6·2	18·8
to end of 20 years	12·13%	3·27	·1	3·3	5·9	18·8

(*Note:* all are net of expenses; taxation is ignored.)

4.14 However, although on average this policy is very successful, it is not guaranteed to be always the best. Table 4.3 shows the number of occasions in each of the 1,000 simulations where the dynamic policy produced the highest proceeds and where investment wholly in shares or in Consols (net of expenses) did. Although the dynamic policy comes top in between 75% and 80% of the 1,000 simulations, shares and Consols share the remaining 20% or more in a roughly 60:40 ratio. Further, the worst results from the dynamic policy are worse than the worst results for Consols, though the worst results for shares are even more disastrous (proceeds of £60 for £100 invested for shares at the worst; £290 for Consols; £190 for the dynamic policy looking ahead 3 years).

Table 4.3. *Number of occasions out of 1,000 simulations where each investment policy produced the highest proceeds*

Looking ahead	*Dynamic policy*	*Shares*	*Consols*	*Cash*
1 Year	759	166	75	0
2 Years	791	143	66	0
3 Years	798	126	76	0
4 Years	784	130	86	0
5 Years	784	132	84	0
to end of 20 years	757	134	109	0

4.15 Before all my readers rush to instruct their stockbrokers to switch their assets wholly into shares or Consols as the case may be depending on current conditions at the time they happen to read this paper, they may like to reflect on the reasons why these apparently glamorous returns cannot be achieved. First, I assume that it is in fact possible to switch without affecting the prices of the securities that are being bought or sold. This is quite unrealistic if large amounts are being switched, and particularly if everybody is wishing to switch in the same direction. The consequence would be that, if all investors in the market made their forecasts on the same basis, prices would rapidly move so that the expected returns from different investments were sufficiently similar for it not to be worth paying the costs of switching. This would have the effect of stabilizing the market, and changing the parameters of my model in an appropriate way.

4.16 I also make the assumption that the true parameters of the market are known to me. I have not investigated what the results would be if I looked ahead using one set of parameters, when the actual outcomes in my simulated results were derived using a different set of parameters. I might prove to be significantly wrong in my forecasting. Finally, even if all my assumptions were correct, there would still be a 20% chance of my dynamic policy not being the best one over the next twenty years. So please don't come back to me in AD 2005 and complain that my dynamic investment policy was not successful for you. If I were to have followed it myself, and it had proved unsuccessful, I should be feeling sore too!

5. INDEX-LINKED ANNUITIES WITH MINIMUM MONEY GUARANTEES

5.1 I described some aspects of what follows in a paper "The Cost of Minimum Money Guarantees on Index-linked Annuities" that appeared in the Transactions of the 22nd International Congress of Actuaries, Volume 2, p. 137. The following description omits some of the details of that earlier paper, but takes further the considerations of how to fix a suitable premium.

5.2 Index-linked stocks are now sufficiently familiar in the United Kingdom not to need description. The concept of an index-linked life annuity is straightforward. A number of companies have already issued such contracts. In this example I shall ignore necessary practical details such as the delay in publication of the Retail Prices Index and the fact that the amount of an annuity payable monthly would probably be revised only annually. Here I assume that all annuities are paid annually in advance. Payment in advance is tidy in order to define the first payment as 1. A straightforward index-linked annuity offers a benefit which is strictly proportional to the Retail Prices Index applicable at the date of payment. However, an index-linked annuity may be written with at least three forms of minimum guaranteed benefit. Let the Retail Prices Index at time t be $Q(t)$, and let the payment at t for an annuity of type i be $Bi(t)$, with $Bi(0) = 1$, for all i; i.e. all the different types start with an initial payment of 1.

5.3 The straightforward annuity is of Type 1, and gives:

$$B1(t) = Q(t)/Q(0).$$

The weakest form of guarantee, Type 2, is that the amount of the annuity will never be less than the initial amount. We thus have:

$$B2(t) = \text{Max}(1, Q(t)/Q(0)).$$

The next form, Type 3, promises that the amount of the annuity will never be less than it was on any previous date, so that it never decreases:

$$\begin{aligned} B3(t) &= \text{Max}\ (Q(u),\ 0 \leqslant u \leqslant t)/Q(0). \\ &= \text{Max}\ (B3\ (t-1),\ Q(t)/Q(0)),\ t \geqslant 1. \end{aligned}$$

The strongest form, Type 4, promises that the increase each year will equal the rate of inflation if that is positive, so that there is a ratchet effect on the amount:

$$B4(t) = B4\ (t-1) \times \text{Max}\ (1,\ Q(t)/Q(t-1)),\ t \geqslant 1.$$

5.4 In all cases: $B1(t) \leqslant B2(t) \leqslant B3(t) \leqslant B4(t)$, and if the Retail Prices Index never falls, then the $Bi(t)$ are equal for all i. If the RPI falls and rises again to a new maximum, the benefits for each of the first three types become the same again, but for Type 4 the ratchet operates, with a permanent increase in the level of benefit. The cost of the guarantee clearly rises with i.

5.5 I can think of no conventional actuarial method that is of any assistance in estimating the cost of these guarantees. A stochastic approach seems essential. However, I shall limit the stochastic element to inflation only, simplifying the investment side by assuming that all investment is in index-linked deposits which yield 3% real per annum at all times. I assume a male aged 65, subject to PA(90) mortality, and make no allowance for stochastic variation in the date of death. (Alternatively, I assume an infinitely large number of such males, who experience PA(90) mortality exactly.)

5.6 I start by simulating inflation using the Full Standard Basis for sufficiently many years. Now that I know what inflation is, in this simulation, I can calculate the actual amount of the benefits that will be paid. I make allowance for the probability of payment, and discount at a real rate of 3%; strictly, since I know what inflation will be I know what the money yield on index-linked stocks will be too. This process gives me a discounted present value for the annuity on the usual lines. For the Type 1 annuity it is always 11·72, which is the value of $\ddot{a}_{65}$ at 3% on

PA(90). For the other Types the present value is sometimes the same as this, the guarantee not having been applicable. Often, however, it is higher. Using my usual 1,000 simulations, I can build up a frequency distribution of the present value of the different types of annuity. Notice how I am treating the present value of the annuity as a random variable.

5.7 The distributions of these present values are extremely skew. Each has a considerable number of values of exactly the minimum amount of 11·72, with a thin tail that stretches upwards from this. Table 5.1 shows the mean and selected higher quantiles of these distributions. It is easily seen that both the mean and the quantiles increase with the strength of the guarantee, and that the extreme values would justify a substantial loading if the office wished to be sure that it had sufficient to meet the payments. At the extreme, to cover a Type 4 annuity with a 99·5% probability would require a consideration of 16·1, some 37% higher than that for a Type 1 annuity.

Table 5.1.
Present values of annuities of various types, Male aged 65

		Quantiles			
Type	*Mean*	*95%*	*97·5%*	*99%*	*99·5%*
1	11·72	11·72	11·72	11·72	11·72
2	11·83	12·3	12·8	13·8	14·2
3	12·01	13·0	13·4	14·2	14·6
4	12·44	14·2	14·7	15·6	16·1

5.8 One could express the loading either as a percentage, as I have just done, or as a constant amount, or as an adjustment to the rate of interest. Experiments show that for different ages at entry the appropriate loading is more nearly a constant percentage. Table 5.2 shows the same results for different ages at entry, expressing however the quantiles in terms of the percentage loading on the minimum (Type 1) values. It can be seen that these decline slowly with increasing age.

Table 5.2.
Present values of annuities of Type 3, Male, various ages

	Type 1	*Percentage loading for quantiles:*			
Age	*Annuity*	*95%*	*97·5%*	*99%*	*99·5%*
65	11·72	11	14	21	25
70	9·79	10	13	20	24
75	7·99	9	13	18	21
80	6·38	8	11	14	19
85	5·02	6	10	12	16
90	3·91	5	8	10	13

5.9 One way of determining what premium to charge would be to choose some probability level and charge an amount corresponding to that probability. The trouble with such an approach is that in a great many cases the office will have overcharged. In a small number of cases the office will have undercharged and will have insufficient funds to meet the annuities that are due for payment. If the 'probability of ruin' for this particular portfolio is to be acceptably small, then the premiums are unacceptably large, and vice versa.

5.10 We have to remember that there is no averaging of experience across different policies which start at the same date. All are subject to the same conditions in respect of the same Retail Prices Index. There may be some slight averaging out over policies written in successive years, but even these are likely to be correlated, and indeed for a Type 4 annuity are wholly correlated in any one year.

5.11 An alternative approach is to say that a life office which writes this type of non-profit business must have capital. Whether this comes from shareholders or from with-profits policyholders is not relevant here. I shall use the term 'shareholders' to include both sources of capital. I assume that the shareholders are required to set up reserves on a very strong basis, for example at the 99·5% probability level. (An even higher probability might be appropriate.) For a male aged 65 this would imply percentage loadings over the basic 3% reserves of 21·1%, 24·5% and 37·3% respectively for the three Types of annuity. I assume that in each future year the office will hold reserves in respect of each surviving annuitant on the same basis, namely a PA(90) 3% reserve with the same loadings as just stated. These annuity values will be applied to the amount of benefit currently in force, which may of course be higher than for a Type 1 annuity if the guarantee is currently operative. While this basis is strictly appropriate for a Type 4 annuity, it is somewhat too high for a Type 2 or Type 3 annuity, because the RPI may have to catch up a bit before the annuity amount actually increases. But I don't mind, at this stage, if the office is over-reserving. In ·5% of cases, 1 in 200, these reserves will prove insufficient anyway.

5.12 Because the reserves are on a stronger basis than the actual experience for most years, there is usually a release of reserves when the experience of a cohort of policies is followed through. This release of reserves is available to shareholders as profit. In at least ·5% of cases the profits will prove to be negative, since the initial reserves will have been insufficient. But in many more cases than this the profit in particular years will prove negative, because profit that has emerged in earlier years has been released. In these circumstances I assume that the shareholders are prepared to provide further capital, in order to increase reserves to the strong level required, since at that level the policies will remain profitable on average. This is an area where perhaps more investigation needs to be done.

5.13 The present value of the profit that will be released in future is a measure of the value of the policy to the shareholders. It therefore represents the amount that the shareholders should be willing to put up in the first place in order to write the policy. This present value of profit differs in each simulation, so it is also a random variable. But it is reasonable to assume that the shareholders have many different independent investments, and are therefore happy to invest their capital on the basis of the average present value of profit, subject to their receiving a sufficient expected return thereon.

5.14 I have not mentioned the rate of interest at which the future profit should be discounted. If it were discounted at 3%, the same rate as earned on the assets, then of course the expected value of profit will be equal to the difference between the initial reserve and the mean present value of the annuities quoted in Table 5.1. But I assume that 3% is the yield obtainable on risk-free index-linked securities. The shareholders can get this without carrying any of the risks inherent in these annuities. What they are being offered is a risky investment, for which they will require a higher rate of return. The rate of return they require for a particular type of risk has to be determined in the long run in the market-place. In this example I have assumed that the shareholders will require a real return of 5%, which is somewhat higher than the 4·5% mean real return of shares inherent in my model, as can be seen from Table 2.2.

5.15 The investment can be thought of as riskier than shares in general, since the distribution of profit is extremely skew, with a fixed upper limit not very much higher than the mean, and a significant probability of its being substantially lower than this or even zero. In a few cases the present value of profit turns out to be negative, and as already noted the shareholders may have to supply additional capital even where the present value overall is positive.

5.16 I have therefore calculated the present value of profit discounted at a real interest rate of 5%. In any year where there is a loss I have actually discounted this also at 5%, which is perhaps inappropriate; I perhaps should have used 3% for discounting losses.

5.17 Table 5.3 shows certain statistics for the present value of shareholders profit, calculated as explained above. It can be seen that in over 1% of cases the discounted profit proves to be negative. Table 5.4 shows how the premium to be charged is calculated, by deducting the shareholders contribution from the initial reserve. These are seen to be somewhat less than the 95% quantiles in Table 5.1. The loadings in excess of the premium for the Type 1 annuity are reasonably competitive. However, the shareholders do require to put up a fair amount of capital in order to write this business.

5.18 In the course of the simulations I counted also the number of years in each simulation in which the profit was negative, and the shareholders would be called upon to make up the difference. Table 5.5 shows the distribution of these

Table 5.3. *Present value of shareholders' profit, Male aged 65*

Type	*Mean*	*Standard deviation*	*Quantiles 1%*	*Quantiles 5%*	*Number of cases negative out of 1,000*
1	·0	—	—	—	—
2	1·99	·36	−·1	1·4	15
3	2·16	·46	·0	1·3	13
4	3·06	·76	·1	1·5	12

Table 5.4. *Premium required from annuitant*

Type	*Initial reserve*	*Shareholders capital*	*Annuitants premium*	*% Loading*
1	11·72	·0	11·72	—
2	14·20	1·99	12·21	4·2
3	14·59	2·16	12·43	6·2
4	16·09	3·06	13·03	11·2

Table 5.5. *Distribution of number of years with negative profit*

Years	*Type 2*	*Type 3*	*Type 4*
0	765	202	328
1	88	216	256
2	52	185	146
3	37	127	112
4	19	111	76
5	12	64	43
6	11	45	17
7	6	20	9
8	5	19	12
9	2	4	1
10	3	5	—
11	—	1	—
12	—	1	—
Total	1,000	1,000	1,000

numbers. Looking at a few individual simulations suggests that calls on the shareholders are more common in the early years of the experience. It is obvious why this is the case for the Type 2 guarantee; after sufficient many years have passed it is very unlikely that the RPI falls below its initial level. For the other types the explanation seems to be that the reserve, calculated on the basis of a constant percentage loading on the 3% reserves, becomes steadily stronger and so less likely to prove insufficient. This can be seen from Table 5.2.

5.19 Clearly there are ways in which my example could be refined. The required reserves could be calculated in a less crude way. And losses could perhaps be dealt with in a different way than I have treated them. There still remain two matters of judgement in applying this model: the 'probability of ruin' on which the strong reserves are based; and the rate of interest required by shareholders in excess of the risk-free rate. However, I believe that focusing the actuary's attention on these questions is helpful. And as I stated above, I see no way that conventional actuarial techniques could assist with this example; the only other possible method of pricing is a purely market-oriented one. Charge what you think you can get away with, and hope not to put off prospective purchasers too much!

6. ALLOWING FOR VARIABILITY IN THE PARAMETERS

6.1 In all the previous simulations, including those quoted in the Faculty paper, I assumed that the parameters of the model were known and fixed. In the Faculty paper and in OARDP 36 I show the effect of varying the parameters, one by one or in groups, in particular ways, to get an impression of the variability of the future returns. Another method of approaching this is to allow the parameters themselves to vary stochastically.

6.2 Thus at the beginning of each simulation I choose the parameters, which are then fixed for that one simulation. I assume that each parameter is normally distributed, independently of the other parameters, and I supply the program with the means and standard deviations of these parameter distributions. The means of my parameter distributions are the same as the fixed parameters in the Full Standard Basis. The standard deviations of parameters that I have chosen are approximately equal to the standard errors of the parameter estimates, based on the statistical investigations described in OARDP 36. The values of all these are shown in Table 6.1.

6.3 The lower part of Table 6.1 shows the fixed initial conditions, which I have held fixed throughout the following simulations, being the same at the start of each simulation, and equal to their overall mean values. However, this means that each simulation is not starting at a neutral position for that particular simulation. It would of course be possible to base the process on initial conditions that were either chosen at random from some specified distribution, or chosen to be the neutral conditions for the start of each simulation. I have not used these alternative methods.

6.4 In order to keep the results of the simulations within reasonable bounds, it is necessary to put certain restraints on the sizes of some of the parameters. For example, standard deviations cannot be negative, nor can mean yields. If the parameters of the autoregressive processes exceed 1, then that process is unstable and likely to be explosive. This is perhaps not theoretically objectionable, but computers do not like numbers bigger than 10 to some very large power, and even a single such number makes nonsense of any calculated means. I have therefore imposed the following limits on parameter values:

QA, *YA* must not exceed 1·0;
QSD, *YSD*, *DSD*, *CSD* must not be less than ·0;
DD and *CD* must not be less than ·0;
YMU and *CMU* must not be less than ·5%;
*CA*1 is restricted so that the total of $CA1+CA2+CA3$ does not exceed 1·0.

This last restraint is slightly arbitrary. A possibly better way of considering the three autoregressive parameters is to factorize the expression:

$$1-CA1.B-CA2.B^2-CA3.B^2,$$

and treat the three resulting coefficients of B as independent parameters, each

Table 6.1. *Values of Mean and Standard Deviations of Parameters*

	Mean	Standard deviation
Inflation:		
QMU	·05	·015
QA	·6	·1
QSD	·05	·005
Share Yield:		
YW	1·35	·35
YMU%	4·0	·5
YA	·6	·1
YSD	·175	·015
Share Dividend:		
DW	·8	·2
DD	·2	·06
DX	·2	·2
DY	−·2	·05
DMU	·0	·02
DB	·375	·15
DSD	·075	·005
Consols Yield:		
CW	1·0	·1
CD	·045	·01
CMU%	3·5	1·0
CY	·06	·05
CA1	1·20	·1
CA2	−·48	·1
CA3	·20	·1
CSD	·14	·015
Initial values (fixed):		
QI	·05	
YI%	4·27932	
YEI	·0	
DEI	·0	
DM	·05	
CI%	8·5	
CM	·05	

with its own mean and standard deviation, and each limited to be not greater than 1·0. However, it was easier to apply the limit in the way that I have done.

6.5 The results of 1,000 simulations for 100 years each are shown in Table 6.2, which should be compared with the results in Table 2.2. The mean values are fairly similar, as they should be. The standard deviations are generally higher, though not as much higher as I had expected. Of course, when translated into final proceeds, rather than mean annual rates of return, the variability becomes fairly large.

6.6 I had expected that the procedure of picking parameters at random from a

Table 6.2. *Results on Full Standard Basis with Stochastic Parameters*

Term (years):	*1*	*5*	*10*	*15*	*20*	*30*	*50*	*75*	*100*
Mean Rate of Inflation (*GQ*):									
E(*GQ*)	5·36	5·27	5·14	5·06	5·05	5·10	5·13	5·14	5·13
SD(*GQ*)	5·56	4·94	4·18	3·75	3·38	2·98	2·44	2·16	2·02
Mean Rate of Money Return on Shares (*GPR*):									
E(*GPR*)	11·64	9·93	9·97	9·98	9·84	9·84	9·91	9·94	9·96
SD(*GPR*)	22·23	8·56	6·18	5·53	5·10	4·70	4·15	3·86	3·71
Correlation Coefficient:									
C(*GPR,GQ*)	−·24	·23	·46	·58	·61	·64	·61	·61	·60
Mean Rate of Money Return on Consols (*GCR*):									
E(*GCR*)	8·54	8·57	8·74	8·84	8·84	8·84	8·92	9·14	9·40
SD(*GCR*)	6·83	3·20	1·80	1·29	1·14	1·34	1·86	2·60	3·59
Correlation Coefficients:									
C(*GCR,GQ*)	−·41	−·60	−·60	−·45	−·12	·23	·50	·45	·37
C(*GCR,GPR*)	·12	−·06	−·22	−·18	·00	·16	·30	·29	·24
Mean Rate of Real Return on Shares (*JPR*):									
E(*JPR*)	6·53	4·57	4·65	4·71	4·57	4·52	4·54	4·57	4·59
SD(*JPR*)	23·38	8·43	5·40	4·34	3·84	3·44	3·11	2·91	2·83
Correlation Coefficient:									
C(*JPR,GQ*)	−·46	−·35	−·27	−·16	−·10	−·03	·00	·03	·03
Mean Rate of Real Return on Consols (*JCR*):									
E(*JCR*)	3·45	3·45	3·63	3·75	3·72	3·64	3·65	3·82	4·07
SD(*JCR*)	10·09	7·13	5·37	4·43	3·66	2·96	2·16	2·42	3·25
Correlation Coefficients:									
C(*JCR,GQ*)	−·81	−·93	−·96	−·96	−·95	−·90	−·71	−·42	−·23
C(*JCR,JPR*)	·39	·36	·28	·18	·12	·04	−·01	·00	·02

distribution before carrying out each simulation would be very complicated to program. In fact it was not. I think I may prefer to use such a fully stochastic model for certain purposes in future. The examples I have given earlier in this paper could each have been based on this model, instead of on the Full Standard Basis with fixed parameters. The investment policy example, as I have implemented it, requires actually two models, which could have been different; one to do the looking ahead, the other to simulate the actual out-turn. One could use a fixed model to look ahead, and then allow the out-turn to be determined by the chosen parameters. Or one could select the parameters for the actual out-turn from their specified distributions, and keep them constant throughout say 100 simulations, repeating the process say 100 times, and recording the frequency with which the dynamic investment policy comes out top within each of the 100 major repeats. One could be even more elaborate and try to use the actual experience of the simulation to estimate its parameters perhaps in a Bayesian way, either taking account of the fact that we actually know the prior distribution of parameters or not. The possibilities are endless.

7. IMPLEMENTATION

7.1 Obviously it is necessary to carry out simulation work of this type with a computer, and some hints on implementing these models may be helpful. Once one has built up a small set of useful subroutines, it is not difficult to put together a program to carry out any particular exercise.

7.2 It is essential to have a convenient subroutine for generating pseudo-random normal variates. Many computer systems have such a subroutine, but some of these, at least in the past, have been less accurate and slower than exact ones. I have used Marsaglia's Polar method which is an exact one, and found it faster than any other that I tried, including the commonly used approximation of adding 12 uniformly distributed random numbers. It is also convenient to take advantage of the structure of binary numbers within a computer, which in the typical IBM 32-bit word machine run from -2^{31} to $+2^{31}-1$, and to take advantage of the fact that, at least in the version of Fortran available to me, overflow of binary numbers is ignored. Thus a linear multiplicative congruential generator written as:

$$IX = 314159269 * IX + 271828189$$

gives a sequence of pseudo-random values of IX within the full range of binary integers.

7.3 The heart of the system is a subroutine to implement the stochastic model. One wants to be able to call it, and ask it to hand back a simulation of N years values, starting with given initial conditions. Besides the obvious values that need to be transmitted, it is useful to have a number of markers also as parameters of the subroutine. One marker could indicate whether this entry for the subroutine is the first, in which case a set of parameter values needs to be obtained, or a subsequent entry, in which case an actual simulation is to be carried out. A second marker may indicate whether to reset the initial conditions to those at the beginning of the first simulation, or whether to continue with the conditions that were left at the end of the last set of years. A third marker could indicate how much, if anything, to print. It is convenient in the development of a program to be able to print out the complete results for a few simulations, showing the progress of every year. But once the program is working one wishes to suppress the printing of individual results, and to show only final tables. However, if the final results show unexpected values it may be desirable to repeat the simulations, printing out say a single line of results for each simulation. It also helps to print the random number seed or seeds at the beginning of each simulation, so that it is possible to start again at any chosen point and print out full details for the simulation that is causing trouble. An expected result may be a genuine, but extreme, consequence of the chosen model. Or it may be caused by a mistake.

7.4 It is convenient to have subroutines available for summarizing the results. I have two main ones. One accumulates the first, second, third and fourth powers of a set of variables $X_1, X_2, \ldots, X_n$ and on request prints out the mean, standard deviation, and skewness and kurtosis coefficients. On request too it accumulates the cross products of the Xs, and gives the correlation coefficients between each pair of Xs. I actually need two versions of this, one to deal with 'real' variables, the other to deal with 'integer' variables; but this is a feature of writing in Fortran, rather than a mathematical requirement.

7.5 The next auxiliary subroutine accumulates a frequency table for each of the variables $X_1, X_2, \ldots, X_n$, counting the number of cases that fall within each cell. From such a frequency table one can pick out the quantiles that may be required. It is necessary to choose a number of cells and a step size beforehand and to have two extreme cells for outliers. If you choose the wrong scale to start with you end up with everything in one of the outlier cells, and have to start again. Another approach, now that computer space seems to be cheap, is simply to record the results for each of the 1,000 or so simulations, and sort them into sequence at the end, thus giving exact values for the kth highest observation. You may notice that the quantiles I quote in the first part of Table 4.1 all end in 5; this is because I am quoting the midpoint of the cell in which the 925th highest and 975th highest of the 1,000 simulations fall.

7.6 With this equipment available one can start on a particular exercise. You should probably start by thinking through carefully what calculations you would carry out given an actual experience in front of you. This shows how to deal with each particular simulation. Then it is desirable to consider what results you wish to record for each simulation, for which you may want to record statistics, frequency distribution, correlation coefficients etc. You may want to record the number of cases that something exceptional, 'ruin', occurs. You may want to record the duration within the simulation at which 'ruin' first occurs, and build up a frequency table of that statistic. It is worth considering these aspects in advance, because elaborate simulations may take a significant amount of computer time, and it is helpful to get all that one wishes out of each run.

7.7 It may be convenient to get out the results for several different exercises within one computer run, for example for different durations of exercise. In my example about Unit Trust expenses, I could have used one set of simulations to give me results for say 5, 10, 15, 20 and 25 years. Actually that is such a simple exercise that it takes very little time to run, and it was probably quicker overall to program for just one duration and run the job a number of times, as I did. Experience and the constraints of each particular machine will show what is best in the circumstances.

7.8 One final hint; I have found it convenient to simplify each problem to its basic essentials. I think it is a mistake to be too realistic in the first place, making careful allowance for mortality, expenses, taxation, etc. when what is interesting is the stochastic variation of investment returns. It is too easy to surround oneself with so many trees that one cannot get out of the wood. Once one understands the essence of the problem, precise details can be added later.