B.A.J. **2**, V, 1057–1193 (1996)

HOW ACTUARIES CAN USE FINANCIAL ECONOMICS

BY A. D. SMITH, B.A.

[Presented to the Institute of Actuaries, 25 March 1996]

ABSTRACT

There has recently been some debate about the usefulness of financial economics to the actuarial profession. The consensus appears to be that the techniques are potentially valuable, but the published material is not quite 'oven ready'. Some development work is required before the material can be applied.

The author has carried out much of the necessary development work on behalf of various clients over the last few years. On some occasions the results have conflicted with more conventional actuarial methods. The main aim of the paper is to bring the new techniques into the public domain so that they can be properly discussed by the profession, and adopted more widely if appropriate.

The paper contains a number of worked examples using techniques developed by financial economists. The author has listed the computer code which generated the examples and deposited copies on the Internet, so that others can explore the issues with a minimal development overhead.

KEYWORDS

Financial Economics; Stochastic Models; Asset-Liability Studies; Dynamic Optimisation; Equilibrium; Arbitrage; Risk-Neutral Law

1. INTRODUCTION

1.1 Financial economics is the application of economic theory to financial markets. Its use in actuarial work continues to be controversial. At one extreme, I have seen models applied uncritically in apparently blind faith, while at the other extreme, some actuaries have claimed that financial economics is part of a conspiracy to displace traditional methods with unsound practices (see, for example, Clarkson in the discussion of Bride & Lomax, 1994).

1.2 Financial economics has produced several specific models which have come into widespread use within the financial community. Examples include the capital asset pricing model, or CAPM, as derived by Sharpe (1964) and the Black & Scholes (1973) formula for option pricing. There are also a number of established models describing the term structure of interest rates and the pricing of associated derivatives, such as Ho & Lee (1986), Hull & White (1990) and Black, Derman & Toy (1990). Several actuarial papers have been published applying such models. Examples include Wilkie (1987) who applies option pricing to bonus policy, or Mehta (1992) who, *inter alia*, demonstrates the use of the CAPM to derive discount rates for calculating appraisal values.

1.3 It is not the purpose of this paper to review established models such as the CAPM or the Black-Scholes formula. Instead, I will examine some underlying techniques which have been fruitful in financial economics. Economists have applied these techniques to classical economic problems in order to derive the established models. In this paper I have applied the same techniques to some current actuarial problems, and obtained solutions which fuse traditional actuarial methods with powerful insights from economics.

1.4 This paper is far from comprehensive. Instead, I have concentrated on a few concepts from economics and applied these in a series of case studies. These are deliberately detailed and 'hands on', since I believe that it is through practical worked examples that we can best see the point of the new techniques. In a higher level description, it is easy to gloss over the various assumptions and approximations which need to be made to get numerical output. Despite the many compromises and heroic assumptions involved, I am still convinced that financial economics can enhance the work of actuaries. Since the paper mainly consists of worked examples, it is necessarily rather piecemeal.

1.5 The intended audience for this paper is actuaries, who are able to appreciate the technicalities of the modelling approaches that I have described. The profession (Nowell *et al.*, 1995) has described financial economics as an area in which we may have to learn some new skills for the future. I have deliberately concentrated on the method, so that the profession may appraise in the discussion whether the material is sound or not. The mathematics is more involved than in many other actuarial fields, and I accept there will be some clients for whom the style of presentation I have adopted is totally unhelpful. Such clients have very good reasons to employ actuaries, since they have little chance of getting the work done in-house. The communication of complex concepts to non-mathematicians is important, but, in my view, ease of communication is never a good reason for using an unsound method. When the profession has reached a consensus on method, that may be the time to start thinking about how it is to be presented to the outside world.

1.6 The remainder of this paper is organised as follows. Section 2 examines the actuarial appraisal value construction, in the light of financial economics. I have considered the use of risk discount rates and locking-in adjustments in a numerical example. Alongside this, I have used the risk neutral methodology developed by financial economists. I demonstrate that these techniques give the same answer, although they come from very different starting premises. Section 3 then reviews the methodology of asset-liability studies from a financial economics perspective. We pay particular attention to the optimisation techniques developed by financial economists, and compare these with the methods actuaries have traditionally employed. Section 4 considers the issues of capital adequacy and risk management, from the framework of maximising shareholder value. Section 5 applies the theories so far developed to build a stochastic asset model. Section 6 considers the issues which arise when we try to ascribe 'market-based' values to entities which are not traded in a liquid market. Section 7 concludes with some general remarks. Appendix A contains some of the more complicated mathematics involved in optimisation. Appendix B contains some notes on calibrating stochastic asset models. Appendix C contains the computer code I have used to generate the examples, which is, in essence, a library of stochastic modelling functions. I have placed demonstration spreadsheets calling these functions on the Internet at www.actuaries.org. I have collected together some of the technical terms in a glossary which forms Appendix D.

2. The Concept of Value

2.1 *The Appraised Value of a PEP*

2.1.1 A large amount has been written on the subject of value, both from an actuarial and a more general economic viewpoint. We can isolate several distinct value concepts. The *present value* is the result of a discounted cash flow calculation. The economic concept of value is the price at which an asset will trade. An *appraised value* refers to the result of a calculation whose output is intended to be consistent with economic value. The *assessed value* or *long-term value* of an asset is an intermediate stage in actuarial pension fund calculations. These are all distinct from the actuarial reserve concept, which is the quantity of assets required today to meet future liabilities (with a predetermined probability). These concepts will differ numerically because of different ways of treating risk, different cash flow expectations or because of market inefficiencies. The purpose of this section is to reconcile three apparently different value concepts, using the example of a personal equity plan, or PEP. The reason for choosing a PEP is its simplicity, in particular with regard to taxation and charging structure. In order to make the issues more comprehensible, I have produced a numerical demonstration rather than algebraic formulae. For typesetting convenience, all contracts are assumed to redeem in five years' time. Some of the relevant algebra is covered in Bride & Lomax (1994).

2.1.2 Daykin has pointed out, in the discussion of Dyson & Exley (1995), that actuaries focus on value while financial economists focus on price. The main actuarial objections to the use of market price is that it is ambiguous, because market values may be volatile over time. It is not so often pointed out that actuarial appraised values are also ambiguous. The ambiguity arises because of the subjective judgement in the choice of basis. In the view of any one actuary, value may be moderately stable in the short term. However, if one talks to several actuaries in succession, one may obtain a series of 'values' which are at least as volatile as market values, if not more so. The key question must be whether a general disregard by actuaries of the information contained in observed prices would have a detrimental impact on the soundness of our financial management. One message of financial economics is that market value is of crucial importance to virtually any financial optimisation problem.

2.1.3 By way of a concrete example, we consider the valuation of a cohort of single premium PEP business, with five years left to maturity. The unit price is assumed to be equal to the market value of assets less selling costs, so that the emerging distributable profits consist entirely of management fees less expenses. The initial situation is as follows:

Initial units in force	10,000
Initial unit price	£100
Annual fee	1%
Expenses per unit	£0.50

2.1.4 Future assumptions, on a best estimate basis, are as follows:

Earned return p.a.	i_{BE}	12.00%
Lapse rate p.a.	q_{BE}	10.00%
Expense inflation p.a.	e_{BE}	3.25%

All policies are assumed to lapse at the end of year 5 if they have not already done so.

2.1.5 We adopt the convention that expenses are incurred at the beginning of the year, while fees are collected at the end. Lapses occur at the year end, after the fees have been collected. We consider only variable expenses, so that fixed costs are added in at a later stage of the calculation. Using the assumptions in ¶2.1.4, we can project the business as follows, up to the 5-year time horizon:

	Policy year				
	1	2	3	4	5
Unit price brought forward	100.00	110.88	122.94	136.32	151.15
Investment return	12.00	13.31	14.75	16.36	18.14
Management fee	(1.12)	(1.24)	(1.38)	(1.53)	(1.69)
Unit price carried forward	110.88	122.94	136.32	151.15	167.60
Units brought forward	10,000	9,000	8,100	7,290	6,561
Lapses	(1,000)	(900)	(810)	(729)	(6,561)
Units carried forward	9,000	8,100	7,290	6,561	0
Fees receivable at year end	11,200	11,177	11,153	11,130	11,107
Benefits payable at year end	(110,880)	(110,649)	(110,419)	(110,190)	(1,099,604)
Expenses per unit	–0.500	–0.516	–0.533	–0.550	–0.568
Expenses payable at year start	(5,000)	(4,646)	(4,318)	(4,012)	(3,728)

2.1.6 We can calculate present values of these quantities using an appropriate discount rate. One possible candidate would be the assumed rate of return on the assets. Projecting these present values into the future, we can then obtain the figures:

	1	2	3	4	5
Present value of fees	40,225	33,852	26,737	18,792	9,917
Present value of expenses	(17,815)	(14,353)	(10,872)	(7,341)	(3,728)

In this example, and for the remainder of this section, the present value is discounted back to the begining of the policy year, so that, for example, the right hand column is the present value of cash flows during policy year 5 as at the start of year 5. The left hand column is the total present value of all cash flows from years 1 to 5 as at the start of year 1. The reader who wishes to see this in more detail will find this worked example in SECT2.XLS on the Internet.

2.1.7 We could also calculate the present value of the policyholder benefits. Subtracting the present value of the fees, we recover the bid value of the units. This makes a lot of sense, as the current assets have to fund both policyholder benefits and also the fees.

2.2 *The Use of Risk Discount Rates*

2.2.1 Actuaries have developed the technique of *appraised value* to describe the value of a block of business to the proprietors of a financial services provider. The underlying techniques are described in Burrows & Whitehead (1987). An essential part of the theory asserts that the discounted cash flow valuation in ¶2.1.5 puts too high an economic value on the business. This can be justified by a comparison between shareholders owning the business itself, or alternatively purchasing the underlying investments directly. Both of these alternatives would provide the same projected rate of return, i.e. the earned rate, on a best estimate basis. However, ownership of the business carries additional risks above those of the underlying investments. Particular risks would include the possibility of heavier than expected lapses or higher than expected expense inflation. The conclusion is that, at the prices in ¶2.1.5, ownership of the underlying investments would be preferable to ownership of the business. In other words, the discounted value in ¶2.1.5 is not directly comparable with the market value basis for the underlying assets. This does not necessarily mean that the present value calculation is pointless, but it does not tell us all that we would like to know about shareholder value.

2.2.2 One means of adjusting for the degree of risk is to postulate a required rate of return on the financial services business. The required rate of return, or *risk discount rate* (henceforth RDR) is used to discount projected distributable profit flows and produce a shareholder value assessment. The risk discount rate exceeds the assumed earned rate by a margin which, in some way, reflects the degree of additional risk assumed. Much can be said about how the price mechanisms in the market lead to certain kinds of risks being rewarded with higher mean returns; consideration of these issues are deferred to Section 6.

2.2.3 There are significant difficulties with discounting at a rate different from the earned rate. If such a technique were to be applied to the underlying assets, then this would give a value below market. This is unfortunate, since the whole point of introducing the technique was to make the appraised value of an enterprise more consistent with the market value of traded securities. For another example, we could consider two PEP providers, A and B, with the same best estimate cash flows. The difference is that B has a greater degree of uncertainty attaching to future expenses. Under conventional methodologies, a higher risk discount rate would be applied when valuing B than when valuing A. This has the intended result of reducing the value of B relative to A. However, the method achieves this implicitly by reducing the present value of future fees, and also reducing the present value of future expenses, but to a smaller degree. Intuitively, we might hope, on the contrary, that an increase in perceived expense risk would result in a larger (i.e. more negative) value being placed on future expenses while the value of fees stays the same.

2.2.4 One solution to the problems in ¶2.2.3 is to apply different risk discount rates to different components of distributable profit according to their different risk characteristics. This has been suggested by Mehta (1992), and I agree that this is theoretically desirable, but it does not yet seem to be widely adopted in actuarial practice, perhaps because of potential difficulties in communicating results to the client.

2.2.5 In our current example, an appropriate risk discount rate for valuing fees would be in excess of the expected earned rate, thus reflecting the additional risk in a lower value. When valuing expenses the reverse applies, since a larger (i.e. more negative) value of expenses is cautious. We, therefore, discount future expenses at a risk discount rate below the earned rate. The parameters I have chosen are as follows:

Fee income RDR	i_q	12.93%
Expense RDR	i_e	8.21%

Certainly, arguments could be made in favour of other parameter values; my arguments for using these are contained in Section 6. On one hand, it could be argued that no risk premium is appropriate in relation to inflation, as inflation is, in some sense, a risk-free asset. On the other hand, one could point out that the real risk is not the underlying inflation in the economy, but rather the possibility of some management action (or inaction) which resulted in unplanned expense overrun. In a sense, the justification for a lower discount rate is not the variability in the cash flows, but an unwinding of over-optimistic mean estimates. The point is that different actuaries may interpret the discount rates differently, but they can all agree on the answer to the calculation. We could argue that for this product the lapse risk is negligible, because, on lapsation, it is only the bid value of units which must be paid. However, it should also be borne in mind that adverse lapse experience will terminate fee income which might have arisen in the future, so that there is at least an opportunity cost to adverse lapse experience. This might justify increasing the discount rate to allow for the risk, as I have done.

2.2.6 Some care must be applied when using risk discount rates to ensure that the quantities being discounted are actually subject to the risks for which the discount rate has been adjusted. In our current example, I have adopted the convention that fees are deducted before lapses are processed, so that the fees in the first year are not subject to lapse risk. Arguably, the fee income should, therefore, be discounted at the earned rate during the year in which it is charged, and at the risk discount rate in prior years. Equivalently, we can fund the fees actuarially on an annual basis. These mean that the fees emerge as a decrease in reserves at the start of the year rather than as a cash flow item at the year end.

Adding in the present value of expenses, we obtain the following appraised values:

	Policy year				
	1	2	3	4	5
Funded value of units	990,000	987,941	985,886	983,835	981,789
Increase in reserves	(10,000)	(9,979)	(9,958)	(9,938)	(9,917)
Present value of earnings	39,644	33,477	26,535	18,720	9,917
Present value of expenses	(18,867)	(15,005)	(11,209)	(7,457)	(3,728)
Appraised value	20,778	18,472	15,326	11,262	6,189

2.2.7 We can solve the equation of value to find the single risk discount rate which, when applied to projected distributed net profit, produces the appraised value above. This provides the following implied discount rates:

	1	2	3	4	5
Net cash flows	5,000	5,333	5,641	5,926	6,189
Implied discount rate	16.70%	16.42%	16.18%	15.97%	

We can explain the decreasing implied discount rate as follows. The effect of lapse experience on distributable profit is geared up by virtue of the subtraction of expenses. As the expense inflation is lower than the assumed investment return, the extent of this gearing falls over time. The falling impact of lapse risk is consistent with a lower risk discount rate.

2.3 *Cost of Capital Adjustments*

2.3.1 There are several schools of thought on how appropriate discount rates may be derived. Traditional approaches would build a discount rate for a project by comparison with the possible returns on alternative projects. Thornton & Wilson (1992) examine historic asset returns in detail, and implicitly assume that the 'best estimate' expected return on assets is an appropriate rate for discounting the liabilities of a pension fund. By the same token, in a demutualisation of a life office, it is sometimes suggested that policyholders would require an earned rate in respect of the profit stream they give up. These approaches are all too often regarded as obvious, and therefore free of assumptions which ought to be tested. On the other hand, other asset pricing models may appear to be the consequence of numerous dubious assumptions, and, as such, are likely to produce suspect results. The abstract construction of a risk discount rate in excess of earned rates may, therefore, be hard to justify, except by noting that a higher risk discount rate adjusts appraised values in the right (i.e. downwards) direction.

2.3.2 An alternative approach is to discount cash flows at an expected earned rate, and to adjust separately for the additional risk outside the discounted cash flow construction. One way to adjust for the risk is to measure the possible deviations from the best estimate projection, and to allocate capital to absorb these possible deviations. The cost of financing this additional capital is then applied as an adjustment to the discounted cash flow, to provide the appraised value. This item has been called the *cost of capital adjustment*, or COCA, by Bride & Lomax (1994). This idea can also be applied to solvency measurement, capital allocation and product pricing, as described in Hooker *et al.* (1996).

2.3.3 The cost of capital adjustment may be applied either as an increase in the liabilities or as a decrease in the value of the assets. Conventionally, the adjustment has often been applied, not to the assets forming the technical reserves, but to those assets constituting the solvency margin. In this case, the COCA is more commonly called a *locking-in adjustment*.

2.3.4 The first step in calculating a COCA is to determine the appropriate amount of capital at risk from adverse fluctuations, which I have called the *sum at risk* (SAR). In the case of lapse risk, all future fees, except those arising at the end of the current year, may be forfeit to heavy lapses. The cost of financing this capital is the risk margin between the required return and that earned on investments. Discounting these costs at the risk discount rate for the fees gives the COCA for the value of fee income:

	Policy year				
	1	2	3	4	5
Sum at risk from lapse	30,225	23,873	16,779	8,855	(0)
Annual capital cost charge	(280)	(221)	(156)	(82)	0
Cost of capital adjustment	(580)	(375)	(202)	(73)	0
Value of fees	39,644	33,477	26,535	18,720	9,917

The last line is the appraised value, obtained by subtracting the COCA from the present value. We could, of course, consider a small amount of capital at risk by using a pessimistic (but less than 100%) lapse assumption, with a larger percentage financing cost, and obtain the same results. Because of this arbitrariness in calculating the sum at risk, I am rather uncomfortable with this approach, and prefer the more direct methods of risk discount rate, as explained above, or risk neutrality, as shown below. I do not believe that regulatory capital requirements for the cohort of business concerned are of much relevance to calculating the sum at risk, and hence the COCA. This is because, in the event of highly adverse experience, capital allocated to other lines of business, or even unallocated retained capital, could be required to meet liabilities.

2.3.5 By the same token, we can evaluate the quantum subject to expense risk as the present value of future expenses, excluding the current year's expenses, which are known. The excess financing cost is again the spread between the required return and the earned rate, which gives the following COCA calculation:

	Policy year				
	1	2	3	4	5
Expense inflation sum at risk	12,815	9,707	6,554	3,329	0
Annual capital cost charge	(486)	(368)	(248)	(126)	0
Cost of capital adjustment	(1,051)	(652)	(337)	(117)	0
Value of expenses	(18,867)	(15,005)	(11,209)	(7,457)	(3,728)

2.3.6 Adding the value of fees to the value of expenses, including the appropriate cost of capital adjustments in each case, we obtain the total appraised value. We notice that the values of fees and expenses are equal to those obtained via the risk discount approach in ¶2.2.6. This is not an artefact of the particular assumptions that I have used, but a general algebraic property. The proof, which is elementary, but tedious, is left to the reader.

2.3.7 I believe that this example can also shed some light on the current debate regarding transfer values in and out of final salary pension schemes. Let us suppose that scheme valuations are carried out on a 'best estimate' basis, with no margins in the assumptions. The actual cash flows required from the sponsor will then be as projected, plus or minus some randomness or noise. Although the noise terms have zero mean, they still represent a cost to the sponsor, as they may be positive at an inconvenient time — for example when cash flow is tight in the underlying business. In order to provide a credible promise of ongoing funding to the members, the sponsor will need to be more heavily capitalised, and exposed to greater risks, than would be the case in the absence of a final salary scheme. This has an economic cost, not funded and off balance sheet, to the sponsor, which may be equated to the COCA. Adding the discounted value of liabilities (at the earned rate) to the COCA, we obtain an appropriate risk adjusted value comparable to a market valuation based on returns from matching gilts, such as might be obtained if the liabilities were to be bought out by an insurance company. There has been a heated debate in the profession as to whether equity or gilt returns are more appropriate for discounting transfer values. Equivalently, this argument comes down to whether the off balance sheet COCA should be transferred together with the assets of the fund.

2.3.8 Let us consider increasing the assumed investment return when valuing a pension fund. This has no direct effect on the actual cash flows, unless the resulting surplus is used to give discretionary increases or other decisions are taken as a consequence of the changed assumptions. However, on paper, the value of liabilities falls. Crucially, the COCA increases by an equal amount, so that the total economic value of the liabilities remains unchanged. Changes in assumed returns are one way of adjusting the pace of funding, or equivalently, of moving the future capital cost onto and off the balance sheet. The COCA is, in effect, acting as a slush fund to smooth out market fluctuations.

2.3.9 Another area where I have applied this theory is in consideration of time and distance reinsurance in Lloyd's and the London Market. This reinsurance provides fixed cash flows in line with an agreed loss projection, rather like a tailored annuity certain. From an insurance perspective, no risk is transferred and the contract is trivial. However, since the cash flows are certain, other features emerge in practical work which would otherwise be obscured by the insurance uncertainty. It is usual for such reinsurance contracts to be matched by the issuer using various sorts of bonds, and for the recoverables to be backed by a letter of credit (LOC). What discount rate should be used when valuing the cash flows: a rate implied from government bonds, some interbank rate, the cost of funds to the reinsured, the reinsurer or the LOC issuer? The discrepancies which arise can be reconciled in terms of cost of capital adjustments.

2.4 *Risk Neutrality*

2.4.1 A further alternative method for allowing for risk in appraisal valuations is to adjust the forecast cash flows. These cash flows are then discounted at a *risk-free* rate, since all the necessary risk adjustment has already been applied to the estimated cash flows themselves. This is equivalent to projecting on a prudent basis, rather than on a best estimate basis. The risk adjustment then emerges as the unwinding of the prudent basis, when actual experience turns out to be more favourable than that originally assumed. This is the idea behind the accruals method of accounting for insurance business. It is also an established technique in option pricing.

2.4.2 Dyson & Exley (1995) move towards this approach. They describe different means of outperforming a gilt benchmark: through superior active management; through investment in corporate bonds; or from the total return on equity-type assets rather than bonds, and conclude that the distinction between these is essentially arbitrary. Any gains from these sources are essentially speculative in nature, and it is therefore imprudent to take advance credit for their success. If they turn out to be successful, then the enterprise concerned will, *ceteris paribus*, outperform the rest of the market. Each investor believes that he or she will achieve this, but not all will succeed, and, in the absence of structural reasons for supposing that particular investors will succeed while others fail, it is not meaningful to crystallise future above-market performance in an appraised value today. The liability side is totally different; benefits from a reduced cost base may well be capitalised into an appraisal value in advance of the emerging cash flow improvements. The distinction is not so much between assets and liabilities, but between a wholesale investment market, which more or less constitutes a level playing field, and a fragmented retail services market, where structural performance margins may persist for considerable periods of time. The distinction is not altogether clear cut; the pricing of investment trusts suggests that the market is assessing the manager's skill *ex ante*, this being factored into the trust price, so that asymmetries in performance are seen by the market as structural to some degree.

2.4.3 I refer to the prudent assumptions as the *risk-neutral* basis. In our example, we consider the following risk-neutral basis:

Risk-free rate p.a.	i_{RN}	7.64%
Risk-neutral lapse rate p.a.	q_{RN}	10.74%
Risk-neutral expense inflation rate p.a.	e_{RN}	3.56%

2.4.4 I propose that the risk-neutral parameters should be calculated according to the following rules:

$$q_{RN} = 1 - \frac{1+i_{BE}}{1+i_q}(1-q_{BE})$$

$$e_{RN} = \left(\frac{1+i_{RN}}{1+i_e}\right)\left(\frac{1+i_q}{1+i_{BE}}\right)(1+e_{BE}) - 1.$$

The motivation behind these relationships is to provide a consistent answer relative to conventional approaches, but we can also see that they make intuitive sense. The risk-neutral lapse rate is adjusted from the best estimate lapse rate by allowing for the difference between the best estimate return and the risk discount rate for fees. The higher the risk discount rate, the greater the margin between the risk-neutral lapse rate and the best estimate rate. Similarly, the risk-neutral expense inflation is adjusted from the best estimate by allowing for the difference between the discount rate for expenses and the risk-free rate. The lower the expense discount rate, the higher the risk-neutral expense inflation. There is also an adjustment for lapse risk in the risk-neutral expense assumption, to reflect the fact that variable expenses are only incurred on units which remain in force.

2.4.5 The business projection on the risk-neutral basis is as follows:

	Policy year				
	1	2	3	4	5
Unit price brought forward	100.0	106.56	113.56	121.01	128.96
Investment return	7.64	8.14	8.68	9.25	9.85
Management fee	(1.08)	(1.15)	(1.22)	(1.30)	(1.39)
Unit price carried forward	106.56	113.56	121.01	128.96	137.42
Units brought forward	10,000	8,926	7,968	7,112	6,348
Lapses	(1,074)	(959)	(856)	(764)	(6,348)
Units carried forward	8,926	7,968	7,112	6,348	0

2.4.6 We can calculate the present value of fees as follows:

Fees at end of year	10,764	10,239	9,739	9,264	8,812
Present value of fees	39,644	31,910	24,109	16,212	8,186

and similarly for expenses:

Expenses per unit	0.500	0.518	0.536	0.555	0.575
Expenses payable at year start	(5,000)	(4,622)	(4,272)	(3,949)	(3,650)
Present value of expenses	(18,867)	(14,926)	(11,092)	(7,340)	(3,650)

which leads to the following appraised value projection:

Appraised value	20,778	16,983	13,017	8,871	4,536

2.4.7 We observe that the initial appraised value is equal to that obtained in ¶2.2.6 and also in ¶2.3.5. However, the projected future appraised values are lower than for the other methods. This is because of the planned margins in the valuation basis. If experience actually follows the best estimate basis, then the risk-neutral appraised value will still agree with the other methods at subsequent points in time. Thus, not only are the appraised values consistent across the three methods, but reported earnings are also consistent.

2.5 *Arbitrage and Value*

2.5.1 The above example is highly simplified. For example, I have ignored tax, fixed expenses, new business, statutory minimum solvency margins and the existence of free assets. Can we be confident that the same principles apply more generally to a full appraised value calculation?

2.5.2 I claim that the answer is *yes*. Let us suppose that we have an algorithm for assigning values to uncertain cash flows. Suppose, further, that this value algorithm has the following properties: firstly, that the value of a sum of cash flows is equal to the sum of the values assessed separately (*linearity*); and secondly, that the value of a sequence of non-negative cash flows is non-negative (*positivity*). Linearity is very important in financial economics, but many actuarial pricing formulae violate it, for example the practice of loadings proportional to standard deviations for pricing reinsurance. If we assume linearity and positivity, then it can be shown, with some additional mild assumptions of a topological nature, that there exists a *risk-neutral* probability law under which all values are obtained by discounting expected cash flows. This risk-neutral probability law has been called the *shadow probability space* by Bride & Lomax (1994), while financial economists tend to call it the *equivalent martingale measure*. It is not usually the same as any assumed 'true' underlying probability law, because the risk-neutral law contains an implicit risk adjustment to mean cash flows. Thus, we can expect that the risk-neutral approach outlined above also applies more generally, so that any internally consistent theory of valuation has a representation in terms of risk-neutral measure. For a proof of the above result, the interested reader is referred to Duffie (1992) or Harrison & Kreps (1979). A good counter-example is the cost of a capital option pricing formula, which, I understand, is due to be published for the first time in Kemp (1996). It may not be immediately obvious that prices calculated according to the CAPM satisfy the linearity condition; in order for this to work, the discount rate will have to change in a particular fashion, as cash flows with different risk characteristics combine. It turns out that these conditions are satisfied by the CAPM, but other more *ad hoc* approaches to risk discount rates may be less satisfactory in this regard.

2.5.3 If market values satisfy the conditions in ¶2.5.2, then it would be impossible to find a series of transactions which required no initial investment, but provided a guaranteed positive return. In other words, there are no opportunities for *arbitrage*. It is still possible for markets to be inefficient, that is, there may be investment classes whose expected outperformance is abnormally high, given the degree of risk. Absence of arbitrage implies that, for such an asset, the chance of under-performance cannot be zero, that is, no trading opportunity is totally risk-free.

2.5.4 We have seen, in Sections 2.2 to 2.4, how three different approaches to risk adjustment turn out to be equivalent. We have shown how a risk adjustment within one paradigm can be transformed into an equivalent risk adjustment within the other paradigms. It is quite valid to produce values either by risk discount rates, or by loading for the cost of capital, or by discounting risk-neutral expectations. Any debate about the merits of the various methods is one of presentation rather than economic substance. There remains one important issue that we have not addressed. How can we determine the appropriate risk adjustment for any of the three methods? This is considered further in Section 6.

2.6 *Examples from Option Pricing*

2.6.1 One important application of the risk-neutral probability approach has been to option pricing. We now examine the dividend-adjusted Black-Scholes formula in the light of a risk-neutral approach. The formula describes the prices at time t of European style options expiring at time u with strike price K on a stock whose current value is S. The stock pays dividends continuously at a yield q, and the rate of interest for a bond of the appropriate term is denoted by r. The dividend-adjusted Black-Scholes formula, due to Garman & Kohlhagen (1983), gives option prices as follows:

$$\text{call} = e^{-q(u-t)}S\Phi(d_1) - e^{-r(u-t)}K\Phi(d_2)$$

$$\text{put} = e^{-r(u-t)}K\Phi(-d_2) - e^{-q(u-t)}S\Phi(-d_1)$$

where:

$$d_1 = \frac{\log\left(\frac{S}{K}\right) + \left(r - q + \frac{\sigma^2}{2}\right)(u-t)}{\sigma\sqrt{u-t}}; \quad d_2 = \frac{\log\left(\frac{S}{K}\right) + \left(r - q - \frac{\sigma^2}{2}\right)(u-t)}{\sigma\sqrt{u-t}} = d_1 - \sigma\sqrt{u-t}.$$

Here Φ denotes the cumulative normal distribution function. The parameter σ is called the *volatility* of the stock. You do not need to understand this formula in order to follow this paper, but its widespread use demonstrates that the applications of risk neutrality are not restricted to trivial special cases.

2.6.2 These results are consistent with discounted expected values, using a discount rate r (continuously compounded). The expected values are calculated on the basis that S_u has a lognormal distribution, given S_t with mean $e^{(r-q)(u-t)}S_t$ and $\text{Var}[\log S_u] = \sigma^2(u-t)$. This is the risk-neutral probability law for the Black-Scholes option pricing model. The total return on the stock from time t to u, allowing for dividend income, is $e^{q(u-t)}S_u/S_t$. We can see that, under the risk-neutral law, the expected value of this quantity is $e^{r(u-t)}$, which is exactly the return on holding a bond. Indeed, under the risk-neutral law the expected returns on all assets are equal, which justifies using a common discount rate in each case.

2.6.3 The risk-neutral probability law does not determine the true probability law. Many true probability laws are possible, and could be consistent with the Black-Scholes formula. However, in order to determine prices, the risk-neutral law is all that is required, and, conversely, the risk-neutral law can be deduced from market prices, provided that enough prices are available. In so-called *incomplete markets*, the necessary prices may not be available, one example being the absence of a liquid market in lapse risk. In such situations, the risk-neutral law concept is still helpful, but there is more than one law consistent with observed prices. The assets for which prices are available are precisely those which have the same value under every risk-neutral law. Within most derivative trading applications, the risk-neutral law is all that is calculated. This is the case with the interest rate models mentioned in ¶1.2. Any true underlying law is more difficult to calibrate, and irrelevant for setting prices.

2.6.4 By contrast, in actuarial applications we cannot usually obtain market quotes for the cost of lapse risk or expense risk, so it is not possible to write down a unique risk-neutral law directly from market prices. It is, therefore, necessary to select the appropriate risk-neutral probabilities by other means. This usually means estimating the true probability law and applying an appropriate risk adjustment. Equivalently, we need to determine the appropriate risk adjusted discount rate. This is considered in Section 6.

2.6.5 It would be possible to price options by discounting expected values according to a real world probability law. However, the appropriate discount rate would depend on the term and on the strike in a somewhat complicated fashion. In the simple case, where the true probability law is a lognormal distribution with a different mean from the risk-neutral distribution, we can plot the appropriate discount rates for call options, as shown in Figure 2.6.5. We notice that very high discount rates are appropriate for options with short maturities and high strikes, where the implied gearing of the option is highest.

2.6.6 Figure 2.6.5 demonstrates that appropriate risk discount rates may be a complicated function of the characteristics of the cash flows. It makes no sense to write down a single risk discount rate to 'adjust for the risk', and subsequently consider alternatives purely on an expected value basis. Rather, if different derivative strategies are to be compared, potentially different discount rates would be appropriate for each strategy. By contrast, the value of each strategy can easily be compared on a consistent basis using the risk-neutral probability law. Thus, for some applications it is easier to use risk neutrality to risk-adjust expected cash flows, for others an equivalent risk adjusted discount rate.

2.6.7 More generally, when several management strategies are to be compared, based on the same underlying uncertainties, a risk-neutral approach is often the best computational algorithm for comparing value. This avoids the need for specific risk adjustments to be carried out individually for each alternative. One application I have seen is the comparison of various vehicles for repaying

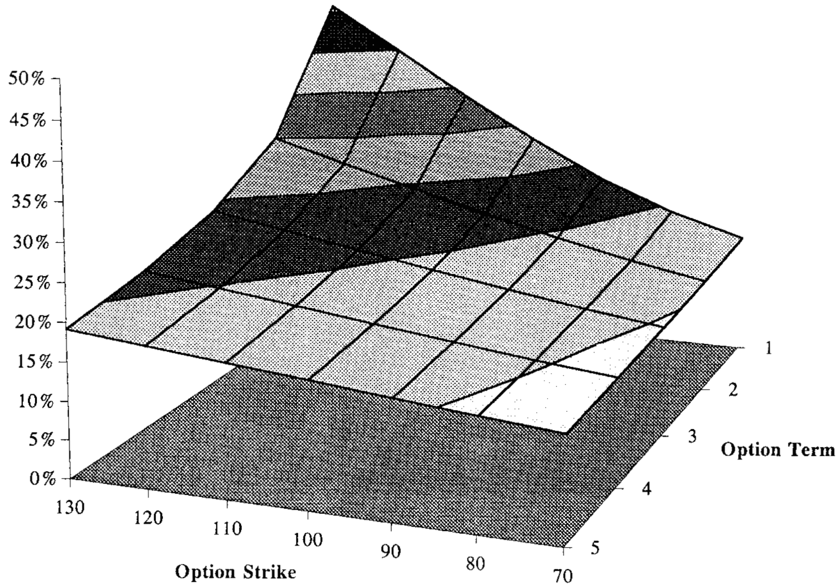

Figure 2.6.5. Discount rate p.a. implied by the Black - Scholes formula

the capital payment on a mortgage, from the consumer's perspective. I examined pensions, PEPs, with-profits and unit-linked endowments, and compared these against the traditional repayment mortgage. Under the risk-neutral probability law, the expected returns are the same for each underlying investment class, so that value for money from the customer's perspective depends principally on issues such as tax and the charging structure. The effect of any higher expected return is unwound by a higher appropriate discount rate, and is, therefore, irrelevant. This contrasts with conventional appraised value techniques, which may attempt to tease out added value by careful manipulations of expected returns and required discount rates. If such calculations were valid, we would have achieved the financial equivalent of a perpetual motion machine, with implications far beyond the appraisal of insurance companies.

3. A Review of Asset-Liability Methodology

3.1 *The Optimisation Problem*

3.1.1 Actuaries have been renowned for their ability to construct models of financial enterprises. It is relatively unusual to see those models designed for any formal optimisation. One exception is the field of asset-liability studies, which is usually viewed within a framework of risk and return. The objective is to maximise the expected return, subject to an acceptable level of risk. Risk can be measured in a variety of ways. One way which has become popular, following

Wise (1984), is to roll up all surplus amounts at an assumed rate of interest to a final time horizon and measure the variance of this rolled-up amount. The set of optimal portfolios which maximise expected returns for a given degree of risk is said to form the *efficient frontier* of *mean-variance efficient portfolios.*

3.1.2 An alternative methodology, following Wise (1987), is to maximise the expected return subject to two constraints, the first being a constraint on the variance, as described above, and the second being a constraint on the initial value of the assets. The technique is sometimes called PEV optimisation, referring to *price*, *expected surplus* and *variance of surplus*. In the case where the liability is a known constant, then it turns out mathematically that PEV efficient portfolios will also be mean-variance efficient. If the liability is not a constant in the normal accounting unit, one can make it constant by defining the liability as the unit of accounting, and, once again, PEV optimisation becomes equivalent to mean-variance optimisation. The extra generality of the PEV approach allows deviations from the mean surplus to be measured in different units from the surplus itself. However, this flexibility is purchased at a price — there are now three dimensions to consider. This is not the limit of the possible complexity. I recently did some work for a property/casualty reinsurer, looking at the efficiency, or otherwise, with which the underwriting book was diversified. There were four variables to trade off, namely: expected profit; variance of profit; premium income and exposure written, so that the efficient frontier is a three-dimensional manifold in hyperspace. The mathematics is an entertaining mix between collective risk theory, modern portfolio theory and geometry. The risk-based capital regime at Lloyd's has now added a fifth dimension to such problems.

3.1.3 Other more complex objectives, together with solutions, are described in Appendix A. While making money for shareholders has not been the primary motivation behind developments in this field, it is interesting to consider the added value to an insurer from an asset-liability study, in economic terms. One way of measuring this is as follows. Suppose that, following an asset-liability study, an investment strategy is determined with an expected return 1% above the previous arrangement, but with the same degree of risk. Then the shareholder value added by the study is the value of an annuity paying 1% of the amount of the fund each year.

3.1.4.1 I am very sceptical of the value added, as calculated in ¶3.1.3, because I think it is too parochial, concentrating only on the insurance fund as an isolated entity. Instead, we ought to look through to the position of a shareholder. This shareholder has already diversified his exposure by holding shares in various sectors, with a sprinkling of gilts, corporate debt and overseas assets. A typical asset-liability study reveals, unsurprisingly, that the total net profit is not efficiently diversified at the insurer level, because of the great concentration in insurance risks. The insurer can potentially increase the expected return for the same level of risk by taking fewer insurance risks, but greater risks to other investment vehicles by way of asset-liability mismatch. This will annoy the shareholder, who has seen diversifiable insurance risk partially replaced by non-diversifiable general market risks. He will then require a higher return, which undoes the effect of the higher mean cash flows on the net present value. Moreover, the insurance share has then effectively become a proxy for the market, so that any tactical reasons for holding it are undermined. It can be argued that the value added to shareholders from the reallocation is zero. This is essentially the argument behind the celebrated Modigliani-Miller theorem (1958), which claims that capital structure is irrelevant when determining the value of the firm. Of course this does not mean that asset-liability studies are useless from all standpoints. One could argue, for example, that directors benefit substantially from having acted on an asset-liability study, since they cannot diversify the company specific risk in the way that shareholders can. Policyholders may also gain from greater security of benefits. When considering a cost benefit analysis for asset-liability studies, it appears that the costs fall on shareholders, but the benefits fall on everybody else!

3.1.4.2 The idea that asset-liability studies do not add shareholder value is unpalatable to many actuaries. One response is to dismiss the Modigliani-Miller theorem, because it fails to reinforce our prejudices or justify our practice. My preferred response is to examine critically the assumptions underlying the Modigliani-Miller theorem, because asset-liability studies could only conceivably add value in a world where these assumptions do not hold. For example, the Modigliani-Miller theorem assumes zero transaction costs. I have relaxed this assumption in Section 4, and, sure enough, capital structure starts to have an impact. A rather different form of asset-liability study can then be justified in terms of minimisation of transaction costs. However, the savings are very much of second order; the value added is nowhere near those implied by the methods of ¶3.1.3.

3.1.5 We now further pursue asset-liability methodology in its conventional form, suspending any scepticism about the value of such investigations until Section 4. It is often found that the optimum for more complex optimisation problems actually lies on, or close to, the mean-variance efficient frontier. Many asset-liability practitioners use this result to construct optimal portfolios by a two-stage process. The first stage is to determine the efficient frontier, and the best choice portfolio is then selected from the efficient frontier, according to the risk appetite of the client. The theoretical explanation of this principle has been explored by Markowitz (1987), while the result has been verified empirically by Booth (1995). Both of these rely, to some degree, on normality or approximate normality of investment returns.

3.1.6 There is a common misconception that mean-variance optimisation should improve investment performance by spotting mispriced assets. However, in order to apply mean-variance techniques, it is necessary to have an underlying stochastic model. The mean-variance optimisation will recommend action on anomalies in the underlying model. The benefits of active trading will only be realised if the underlying model accurately describes trading opportunities. Many models have been built with this in mind; but the track record of actively managed funds in the United Kingdom, based on such models, is not impressive. There is still an *idiot savant* element to the really successful investment managers. Of course, it would be astonishing if mean-variance analysis applied to a simple asset model were, by some accident, to be a good algorithm for stock selection.

3.2 *Components of a Model*

3.2.1 There is more to a full economic model than a stochastic description of asset prices. A complete model should have something to say about four aspects of the economy. Most published models have some pieces missing from the documentation. However, these pieces can often be inferred by close investigation of the model, as we shall see. The aspects (which conveniently all begin with 'p') are as follows:

3.2.2 *Probability*. This aspect constructs joint probability distribution of all cash flows relevant to an enterprise, including both macro and micro economic features.

3.2.3 *Pricing*. This aspect answers the question of how markets will price specified cash flows. One would wish to consider, not only the cash flows arising from underlying 'traded' economic activities, but also the way in which the market would price derivative transactions, or cash flows which are not traded separately, such as the split of a company value by line of business.

3.2.4 *Preference*. This aspect describes how market agents decide which cash flows they like best, and which they like less. A simple approach is to specify a utility function. Actuaries may be more accustomed to optimising the trade-off between some definition of risk and return, taking into account liabilities as well as assets. Each market participant has a view on the market. The portfolio held reflects this view, and also an assessment of the trade-off between risk and return. If we can quantify the probability view and the preferences of the agent, we can repeat his own calculations to infer the likely asset allocation. If this does not reconcile with the assets actually held, then we have either misrepresented the probability view or the preferences of the agent concerned. Thus, these first three aspects determine the *demand* for different assets.

3.2.5 *Prevalence*. This aspect considers the type of assets available in the market for investment, and the capitalisation of each asset class moving forward in time. This is a crude description of the supply side. The total market capitalisation of all asset classes is the sum total of the asset allocation strategies of individual agents. We can consider a representative agent, whose probability assessment and preferences are market averages. The asset allocation held by such an agent will be an average of those held by each agent, which is proportional to the market capitalisation for each asset class. This provides a relationship between the four aspects of a model. It would be possible to go into more detail on the supply side, analysing capital raising in the same depth as investment decisions. I have not done this.

3.2.6 Traditionally, actuaries have been concerned mainly with the probability aspects. The reason for taking a broader view is to make the best use of the data available. Some models explain past market moves; others explain current investor behaviour. By combining the two, one can obtain substantial synergy benefits, and the whole calibration process becomes much more stable. Conversely, I would regard with some scepticism any model built on statistical analysis which does not, in the main, explain current investor behaviour, particularly if the model is to be used to compare investment strategies.

3.3 *A Variety of Asset Models*

3.3.1 There is a variety of different models which have been proposed for asset-liability studies. I have also come across a number of models which are not in the public domain, and obtained specifications with varying levels of detail. It is sometimes claimed that these proprietary models contain commercially valuable insights not available from published material. While I cannot rule this out entirely, the cases which I have seen do not support this view. Instead, the main obstacle to publication has been the amount of work involved in tying up loose ends, defending assumptions, testing hypotheses and documenting results. Wilkie (1995) has set an awesome precedent in this regard. Six types of models, which I believe cover most of the approaches currently in use, are outlined in ¶¶3.3.2 to 3.3.7. The reader is encouraged to experiment with the implementations of these models, which I have posted on the Internet. The code is also listed in Appendix C.

3.3.2 The first type of model that we consider is the random walk model. Such models draw support from the efficient market hypothesis (EMH), which, in one form, states that there is no information in historic price movements which enables a trader to predict future movements profitably. One interpretation of this is to assume that future returns are statistically independent of past returns, with a constant distribution over time. Of course, neither of these are strictly implied by the EMH; for example, it is quite possible, although slightly less convenient, that volatility may show trends over time. I have taken the easier route, and have also assumed that total returns have a lognormal distribution. Correlations are possible between different asset classes, but not between different time intervals. One feature of this model is that the prospective expected return for any single asset class does not change over time. This is hard to defend if returns are measured in nominal terms, but more plausible when adjusted for inflation. The inflation adjusted random walk model in this paper is adapted, with permission, from an unpublished model which M.H.D. Kemp has built. I have recalibrated the model for this paper; more details of the calibration can be found in Appendix B.

3.3.3 At the opposite end of the spectrum are the chaotic models. These assume that the state of the world changes according to some highly complex, but deterministic, laws. These laws exhibit extreme sensitivity to initial conditions, so that, if the initial state is observed within a certain error tolerance, then forecasts become increasingly uncertain as the time horizon moves into the future. This is achieved without any randomness at all, except in setting the initial conditions. I have adapted a chaotic model described by Clark (1992).

3.3.4 There has been some criticism of the lognormal distribution applied to asset returns, largely because it is not well supported by data. Investment data, typically, show a higher frequency of very small changes and of extreme results than a fitted lognormal distribution. Moderate changes are correspondingly over-represented by the lognormal model. One attempt to address this in statistical modelling has been to use so-called stable distributions. This family, discussed in Walter (1990), has the property that the sum of two independent stable variables still has a stable distribution, possibly after a linear transformation. This property is convenient for modelling log returns, since the log return over a long period is the sum of independent log returns over shorter periods. When using stable distributions for log returns, the return distribution over five years has the same form as the return distribution over five minutes. The self-similarity feature is sometimes called the *fractal property*, and models using stable distributions are described as *fractal models*. The models include the normal random walk as a special case; apart from this case, most of the fractal models have infinite mean returns.

3.3.5 There are several features of the markets which are important over the long term, but seem hard to detect from short-term data. For example, some economic series such as the rate of inflation, bond yields, equity dividend yields and total returns may have a long-term mean value. The distribution of these quantities is then said to be *stationary* over time, which means that, over a sufficiently long time period, the observed frequencies of different values will approximate to some long-term distribution. Such a hypothesis underlies the actuarial concept of long-term expected return, although the evidence for stationarity in historic returns is, at best, ambiguous. Stationarity can be implemented using *autoregressive models*, that is models where future movements tend to revert to a long-term mean. Autoregressive models are sometimes described, synonymously, as mean reverting or error correcting. Another observed feature of the market is that, for many asset classes, the income stream is predictable in the short term, even when the capital value is volatile. This obviously applies to gilts, but also, in some degree, to equity and property. There is a case for modelling income and yields separately, producing prices as a ratio of the two. This approach has been adopted by Wilkie (1995). In this paper I have used a subset of the Wilkie model, but I have used the parameters for the property series as suggested in Daykin & Hey (1990) rather than Wilkie (1995).

3.3.6 We have, so far, examined models where returns on any asset class are stationary. We could consider, more generally, those models where asset returns are *cointegrated*, which means that, while return distributions may change over time, there are relationships between returns on different asset classes which do not change. One example is the model of Dyson & Exley (1995), which applies the *rational expectation hypothesis*. They argue that the current term structure of interest rates implies forecasts of all short-term rates in the future. If these forecasts are unbiased, then successive forecasts of the same short rate must perform a random walk over time. The volatility of capital values is due to changes in estimates of the longer-term future cash flows and appropriate discount rates. The model enables future term structures to evolve from the current one, rather than forcing the yield curve to arise from a given parametric family. Similar arguments are applied to inflation expectations implied from index-linked gilts, and also to expected dividend growth assumptions. This results in a model which, like the Wilkie model, produces income streams which are predictable in the short term. However, in the longer term, cash flows are much less predictable for cointegrated models. The long-term unpredictability arises because expected returns are not stationary, but instead, in a cointegrated model, perform random walks in their own right. One implication is that there is no such thing as the long-term rate of return; the average return on an asset class over any period is a random variable, and this random variable does not converge to any limit over long periods. This introduces a dimension of financial risk which is not captured by conventional models, even if these models are stochastic. For example, when simulating the Wilkie model 1,000 times, I obtain 1,000 economic scenarios, each of which, over a sufficiently long time horizon, has 4.7% p.a. inflation. Needless to say, for many actuarial applications, the stochastic model gives results which are not terribly different from assuming inflation of 4.7% p.a. on a deterministic basis. This leads to a common complaint of stationary stochastic models, that all that seems to matter in the long term is the long-term mean return, so there is little insight to be had from stochastic projections relative to deterministic ones. This complaint arises from the stationarity assumption, and does not apply to the cointegrated models. Incidentally, the problem can be overcome, even in apparently stationary models, by choosing the underlying parameters stochastically. As an aside, I note that the rationale for many actuarial methods relies on the concept of a long-term return. If there is no long-term return, it does not necessarily follow that the methods are worthless, but the profession may need to rethink the reasons why the methods work. Dyson & Exley (1995) are concerned mainly with the short-term changes in values (although the term of the cash flows being discounted may be very long). They constructed a model by examining only first order terms. In order to simulate over longer time periods, one must specify the higher order terms, and there is some arbitrariness in this selection. I have taken these higher terms to be zero, which is the simplest way to proceed. For the purposes of the current paper, I have recalibrated the Dyson & Exley model and extended it to produce a series for property.

3.3.7 I construct my stochastic asset model from the four components listed in Section 3.2. More of the detail is explained in Section 5. I describe the model as a 'jump equilibrium' model, since these are its two most distinctive features. I would not claim that my way is the only sound way to build a model; I have worked through the details to show that it can be done, so that a more economic approach can at least be considered as a viable alternative to the currently dominant statistically-based methodologies. Whether it is worth the effort is another issue, which I hope that the profession can now discuss from a more informed perspective, having seen a working example. The construction of such a model is itself an exercise in financial economics. Although the derivation of my model is rather different from that of Dyson & Exley, the resulting models are strikingly similar, and both show a cointegrated structure. The main differences are that I have used gamma distributions instead of normal ones, and I have selected non-zero higher order terms. In addition, I have built my model within an economic framework which allows the consideration of various types of derivatives, which allows me to broaden the scope of the study.

3.4 *The Risk-Return Diagram*

3.4.1 One helpful diagram often shown in asset-liability studies is a *risk-return* plot. This is a scatter graph showing the risk (usually measured as a standard deviation) on the horizontal axis, with the mean return on the vertical axis. Each mix of assets considered will give rise to a single point on the risk-return plot.

3.4.2 The risk-return plot provides a useful algorithm for estimating the efficient frontier. In many situations it can be shown that the feasible region of the risk-return plane is convex. The efficient frontier will then be a concave function which lies above all the sample points. The algorithm I use calculates the smallest concave function which lies above all the points on the risk-return plot. Of course, this will, in general, understate the efficient frontier by a small margin, since it is unlikely that the efficient frontier exactly includes all the asset mixes on my computed frontier. In the examples below I have used 495 sample portfolios to estimate the frontier, and this seems to be sufficiently accurate for most purposes. It is worth mentioning that, in the case of a random walk model, the mean and variance of a constant mix portfolio can be determined analytically from the single period means and variances. In such cases, the efficient frontier for a many-year problem will consist of portfolios which are mean-variance efficient over a single year (or other accounting period). Such portfolios may be determined exactly by means of a quadratic programming algorithm, as explained in Lockyer (1990). Sadly, this result does not appear to extend to models outside the random walk framework.

3.4.3 It is of interest to consider portfolios containing only one or two asset classes. The single asset portfolios will be isolated points on the risk return plot. The set of portfolios containing only two asset classes then forms a curve between the singletons representing each asset class alone. If the two asset classes are not highly correlated, then there is a benefit to diversification, and so some combinations of the two asset classes may have lower risk than either class alone. In this case, the curve on the risk-return plot will be C-shaped. For more highly correlated asset pairs, the curve of mixed portfolios will be less curved and more like a straight line. A 'skeleton' plot of the portfolios containing only one or two asset classes can enable the essential features of the model to be judged at a glance. I also like to add dots corresponding to portfolios with more than two asset components. Figures 3.4.4.1, 3.4.4.2 and 3.4.4.4-3.4.4.6 are examples of such skeleton plots.

3.4.4 We now consider each of these models in turn, and examine the risk-return plots. These show the standard deviation and mean of the accumulation of 1 p.a. in real terms over 5 years. This example has been considered in Booth (1995), in the context of defined contribution pension plans. The efficient frontier is the boundary to the top left of the feasible regions.

3.4.4.1 I have taken the random walk model as my 'base case' model, because readers are most likely to be familiar with it. The risk-return plot is shown in Figure 3.4.4.1. This chart was obtained analytically, in contrast to the other figures in this section, where I used 1000 simulations.

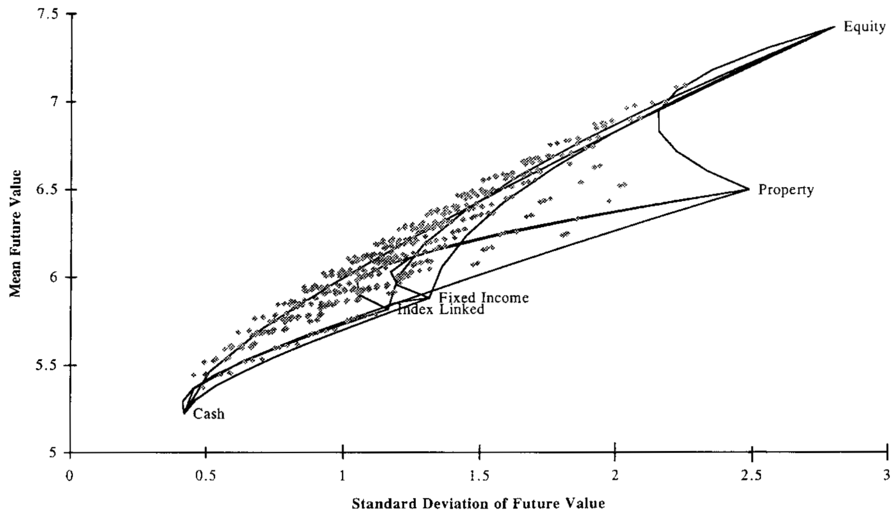

Figure 3.4.4.1. Risk - return plot for the random walk model

3.4.4.2 The corresponding chart for the chaotic model is shown in Figure 3.4.4.2. The major interesting feature is its remarkable similarity to the random walk model in Figure 3.4.4.1. As the chaotic model was easier to program and also runs faster, it could be argued that the chaotic approach should normally be used in preference to random walks. I believe that this conclusion is totally wrong; there are very many important aspects of chaotic models, in particular their short-term predictability, which are not shared by random walks. The apparent similarity is an unfortunate artefact of the way in which we have become accustomed to performing asset-liability studies.

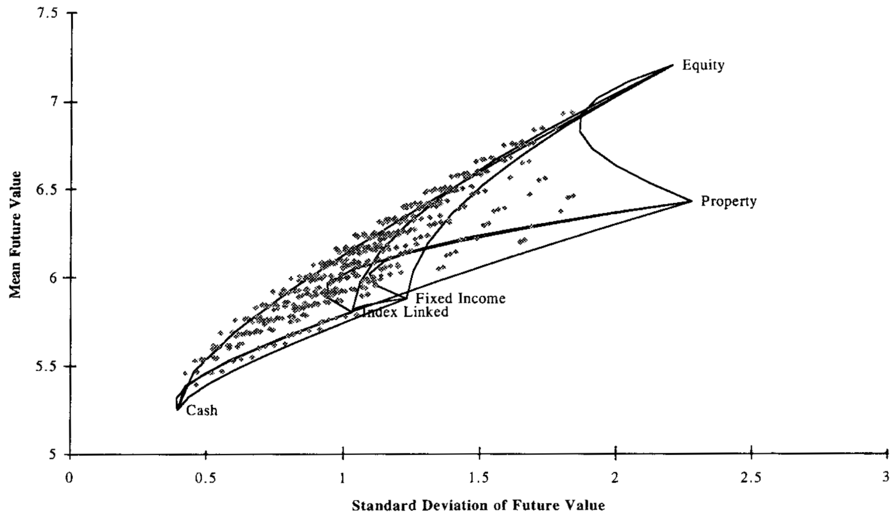

Figure 3.4.4.2. Risk - return plot for the chaotic model

3.4.4.3 A considerable complication arises when considering fractal models in the context of mean-variance analysis — the mean and variance of any asset return are both infinite. However, we can create strategies with finite means and variances by a simple device which I call the *charitable strategy*. Under this strategy, if the asset value on expiry exceeds a predetermined upper limit, then any excess is donated to charity. The commercial merits of such a scheme are dubious, but at least the outcome, being bounded between zero and the chosen limit, has finite mean and variance. We can then plot these strategies on a risk-return diagram, but here each asset mix is a curve (corresponding to different upper limits) rather than a point. It is not easy to plot an analogy of the skeleton plot I have used elsewhere, so I have, instead, been content to plot points corresponding to various asset mixes and upper limits. The results are shown in Figure 3.4.4.3.

3.4.4.4 The skeleton plot for the Wilkie model is shown in Figure 3.4.4.4. Several differences are noticeable relative to the random walk model. The overall level of volatility is rather lower for most asset classes (the most notable exception being cash). There are two causes of this. Firstly, since the Wilkie model has more parameters than the random walk model, more of the historic data are explained by the fitted model, and so the residual noise, which determines the error terms, and hence the volatility, is lower. Secondly, the mean reverting features of the Wilkie model mean that a high real return one year is likely to be followed by a low one the following year, so that the volatility of five-year returns is less than might at first be supposed, given one-year

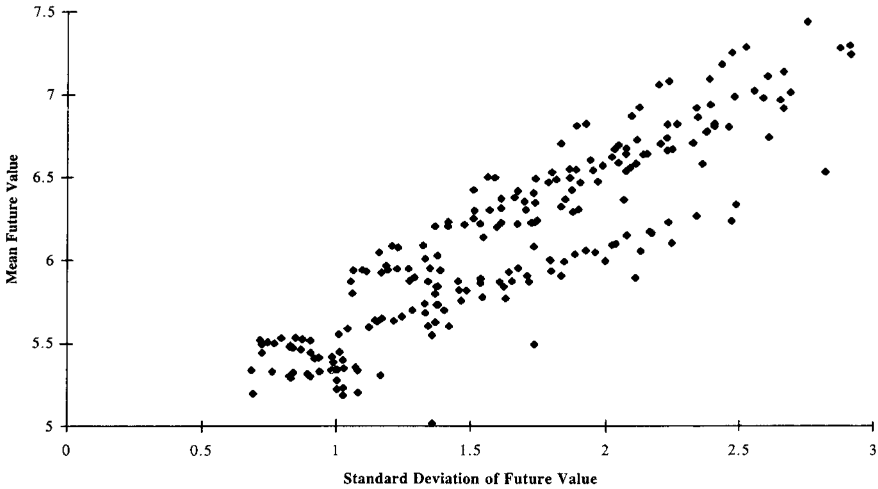

Figure 3.4.4.3. Risk - return plot for the fractal model

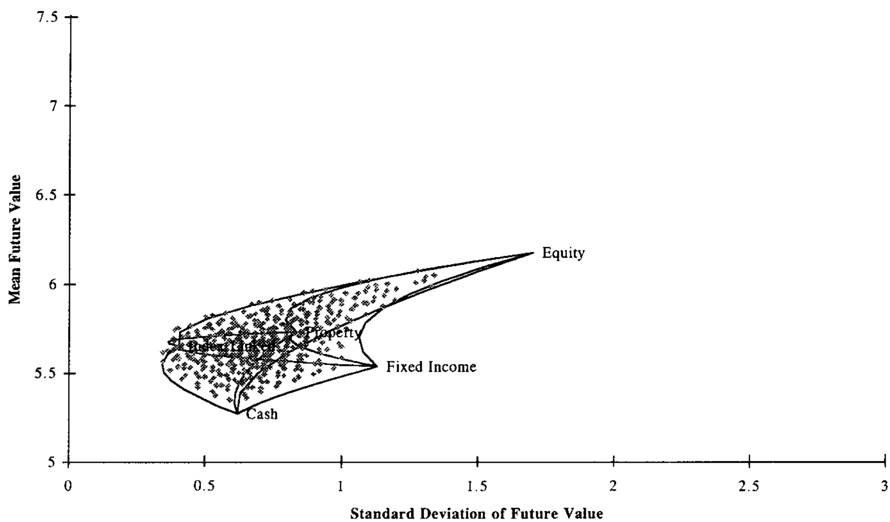

Figure 3.4.4.4. Risk - return plot for the Wilkie model

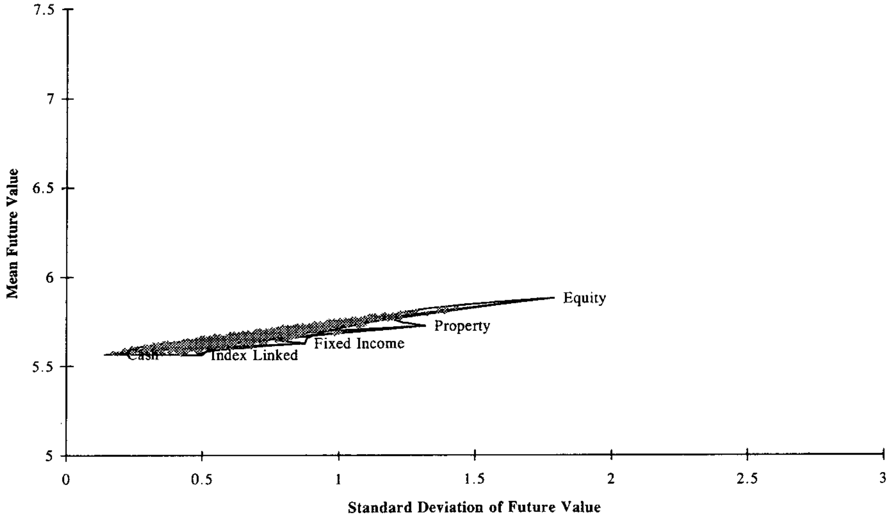

Figure 3.4.4.5. Risk - return plot for the Dyson & Exley model

volatilities. The numerical discrepancy between volatilities is particularly marked for index-linked gilts, where, in real life, the real return over the term to maturity must average out to the gross redemption yield, as occurs in the Wilkie model. Any model for a single gilt which supposes that real returns are independent from year to year, as the random walk does, must, therefore, be suspect (although the assumption is more defensible when applied to a constant maturity index of gilts). The reason for the high volatility of real returns on cash in the Wilkie model is somewhat harder to explain; I make a few observations to shed some light on the matter. We start by noticing that Kemp's version of the random walk model borrows the inflation component from the Wilkie model, so the difference cannot be due to differences in inflation. When examining nominal returns on cash, we see that the reverse effect holds; that is that the Wilkie model produces low volatilities, while the random walk model produces a much more volatile series, even allowing negative returns from time to time. In all fairness to Kemp, his model was designed primarily for pension fund work, where cash holdings are generally small, and where held, it is for reasons of liquidity rather than any more fundamental investment motivation, so that the odd behaviour of cash returns is of little practical consequence. However, notwithstanding the occasional negative interest rates, the volatility of real cash returns in the random walk model is much closer to historic volatilities than the Wilkie model is. One possible explanation of the puzzle is that the inflation component (common to both models) is too volatile over the short term. The random walk compensates for this by producing volatile nominal cash returns, so that, by a cancellation of errors, the volatility of

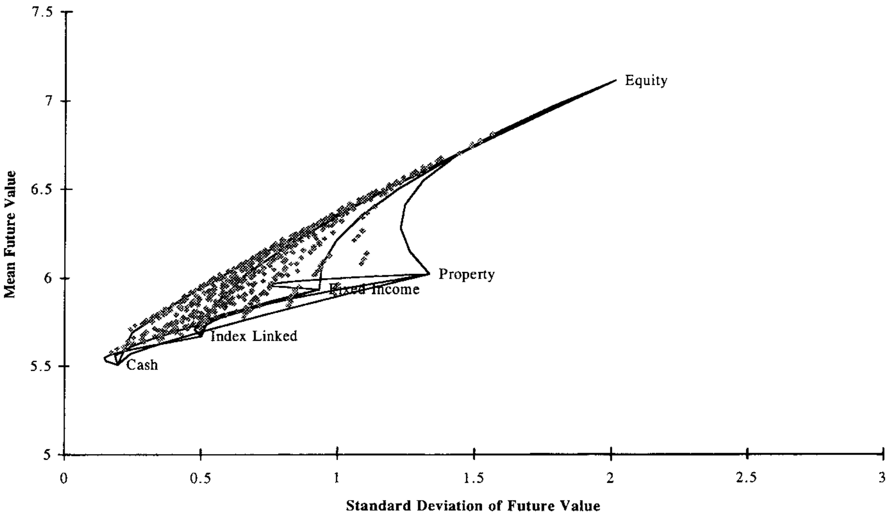

Figure 3.4.4.6. Risk - return plot for the jump equilibrium model

real cash returns is about right. On the other hand, the Wilkie model has a cash model which is appropriately smooth in nominal terms, so that, when expressed in real terms, cash returns appear more volatile than they ought.

3.4.4.5 We can now consider the cointegrated models. The risk-return plot for the Dyson & Exley model is shown in Figure 3.4.4.5. The most striking thing about this chart is that the expected returns from risky assets are much lower than for other models. Indeed, the expected log returns are the same for each asset class. This is largely due to my choice of setting second order terms equal to zero. The volatilities of the Dyson & Exley model are broadly consistent with the random walk model; however, it is interesting to note that cash is significantly less volatile than with the other models. The nominal returns on cash behave similarly for the Wilkie and for the Dyson & Exley models; differences in real cash volatility are a consequence of the different ways in which the models construct inflation. The random walk model (and also the Wilkie model) uses an inflation component with a long-term mean rate of growth. Uncertainty about future inflation is broadly the cumulative total of annual deviations from that mean. By contrast, the Dyson & Exley model produces inflation series which are smooth on an annual basis; uncertainty over the longer term is due to changes in the underlying mean. Thus, if we consider the standard error of the rate of inflation over the next year, the Wilkie model produces a standard error of over 4%, while my calibration of the Dyson & Exley has around 0.7% (their original calibration gives a slightly higher figure). Over the longer term, as observed in ¶3.3.6, these figures swap over, and the Dyson & Exley model has much greater uncertainty. Interestingly, the average error in the Bank of England's one-year-ahead inflation forecasts since 1985 (when they were first made public) has been just under 0.7%, which lends some support to the Dyson & Exley approach; however, the forecasts would probably have been less accurate during other periods, particularly the oil shock of the early seventies.

3.4.4.6 Finally, we come on to the jump equilibrium model, whose risk-return plot is shown in Figure 3.4.4.6. The volatilities are broadly consistent with the Dyson & Exley model, but the expected returns are closer to the random walk approach.

3.5 *Dynamic Optimisation*

3.5.1 We have, so far, considered asset strategies with a constant mix between asset classes over time — sometimes called *static strategies*. It is also interesting to examine strategies where asset allocations can change over time. Such strategies are sometimes called *dynamic*. A dynamic strategy is specified, not by a fixed proportion, but by a series of functions, as follows:

the asset mix between t=0 and t=1 is in specified proportions;

the asset mix between t=1 and t=2 is a function of the economic situation at t=1;

the asset mix between t=2 and t=3 is a function of the economic situations at t=1 and t=2;

and so on ... .

I must emphasise that dynamic optimisation cannot allow looking ahead, so that, for example, it is not permitted to decide on an asset mix between t=1 and t=2 knowing the economic situation at t=2. It is this information constraint which makes dynamic optimisation complex, and interesting.

3.5.2 There are several business reasons for considering dynamic optimisation, as outlined below.

3.5.2.1 One reason for considering dynamic strategies is to allow continued readjustment of a matching portfolio. For example, one may wish to examine the effectiveness of a delta hedging strategy for replicating an option or for portfolio insurance.

3.5.2.2 Another reason for examining dynamic strategies may be to exploit asset classes which are temporarily cheap or dear. Thus, there is a potential for quantifying the benefits of an active investment management approach.

3.5.2.3 A further reason for examining dynamic strategies is to allow the risk-return characteristics of the asset portfolio to be adjusted as a function of solvency levels.

3.5.2.4 Risk averse investors may wish to reduce exposure to risky assets as markets fall, and gear up exposure on a rising market.

3.5.3 When we considered static strategies, we were able to identify a finite number of asset mixes such that any possible static strategy comes close to one of the specified finite selection. This facilitated the estimation of efficient frontiers from a risk-return diagram. On the other hand, there are a huge number of possible dynamic strategies, since future investments can depend on the current economic environment in any way whatsoever. This means that it is difficult to come up with a finite list of representative dynamic strategies in such a way that any strategy is reasonably close to one in the list. The kind of trial and error approach which I employed to optimise a static strategy no longer works for dynamic optimisation.

3.5.4 Since static strategies form a subset of dynamic strategies, we would generally expect the dynamic optimum to improve on the static optimum. Nevertheless, it is rare to find dynamic optimisation employed in actuarial work. I have asked some clients why this is the case, and three reasons came out.

3.5.4.1 Dynamic optimisation is perceived to be too difficult. It is certainly true that dynamic optimisation is harder than static optimisation. It would be nice to have a generic 'sledgehammer' approach for attacking this sort of problem. For example, we might attempt to use simulations. Let us suppose that we have a ten-year time horizon and that one thousand simulated outcomes are required in order to make a decision. When making a decision at year 1, we will need to take into account the scope for changing those decisions later in the process. Thus, for each simulation starting at time 0, we need to solve an optimisation problem to solve the optimal strategy from time 1 onwards, each of which may require 1000 simulations. The argument continues — in order to take account of the decision process at time 2 we will need 1000^3 simulations in all. For the total problem we are looking at 1000^{10} simulations. If we manage 1000 simulations in a second, the total solution will take far longer than current estimates of the age of the universe. Some faster approaches to the problem of dynamic optimisation do exist, based on approximating the model with a finite decision tree. However, these algorithms converge very slowly when return distributions are not normal or when many asset classes are involved. For the kind of problems we consider here, and at the current state of computer technology, the sledgehammer approach is impractical. However, dynamic optimisation is not so difficult as to be impossible, and I have developed some *ad hoc* methods to optimise each of the six models discussed in Section 3.3. The results are plotted below, and the methods are outlined in Appendix A.

3.5.4.2 It is sometimes claimed that there is little to be gained by the use of dynamic optimisation, and that static or near-static strategies are likely to be optimal anyway. Intuitively, I feel that this ought to be the case for many actuarial problems. However, when I tested this hypothesis I found dynamic trading was advantageous in every model except the fractal model.

3.5.4.3 Dynamic optimisation is stretching the limits of what we might reasonably expect models to be able to do. All models will have their flaws, and dynamic optimisation has a habit of homing in on the particular problems of the model in hand. Optimisation may suggest highly geared investment strategies, that is, short-selling large quantities of an under-performing asset and using the proceeds to purchase out-performing assets. Such strategies may appear highly profitable using one particular model, but provide the opportunity to lose large sums of money if the model is wrong, and should not, therefore, be implemented recklessly. I have some sympathy with this view, but I do not see why the solution is to restrict attention to static strategies, as is commonly done. Instead, a better way of allowing for model risk is to check any recommended strategy on several models. The restriction to static strategies results in two sources of error — firstly a vulnerable model, and secondly an optimisation which is artificially constrained not to find the optimum. If these two errors cancel out it will be purely fortuitous. Furthermore, not all models are equally vulnerable under dynamic optimisation, as my empirical results below show.

3.5.5 I have approached the dynamic optimisation from two sides. Firstly, I have used some informed guesswork to construct a series of strategies that may be nearly optimal. Secondly, I have determined some theoretical upper bounds on what can be achieved, although these upper bounds are non-constructive, in that they do not deliver a strategy by which the bounds can be attained. Where the results of my guesswork are close to the theoretical maximum, I can be confident that my method is nearly optimal, and furthermore that the theoretical inequalities are reasonably tight.

3.5.5.1 The dynamic strategies that I have tried are as follows. I start by setting a target final amount (in real terms). At each point in time, I perform a mini asset-liability study with a one-year time horizon, trying to match the asset share at the end of the year to a special purpose 'reserve' figure. The reserve is the present value of the target minus the present value of future premiums, using a market- based real discount rate, similar to the principle behind a bonus reserve valuation (and bearing very little relation to any statutory reserve calculation). I perform the mini-study using a small number of scenarios, because of run-time constraints. These scenarios are not chosen randomly, but rather deterministically, using points selected as in a numerical integration exercise. I have used 6 points for the cointegrated models, 8 points for the random walk and variants, and 10 points for the Wilkie model. Different targets produce various points on the efficient frontier. These choices of dynamic strategy are based on a hunch that they might improve on the static method, so we should not expect them to work well for all models. I have not demonstrated that they are optimal in any sense; merely that they are sometimes better than what we had before.

3.5.5.2 The upper bound is obtained using the dual Fenchel conjugate construction, which is described in Appendix A.

3.5.6 We can judge the effectiveness of a dynamic investment strategy by looking at how far we have been able to advance the efficient frontier to the top left. A good strategy will roll back the frontiers to a greater extent than a poor strategy. An example of a poor strategy is to use a passive reserving basis in conjunction with assets valued at market. As the volatility of the liabilities will appear low, this suggests investment in short-dated bonds. The long-term impact of such a strategy is a duration mismatch, which results in substantial risk over a five-year time horizon, with little compensating reward. The choice of reserving basis is clearly key in formulating the short-term strategy. We can regard a basis as being 'useful' if action, using a series of one-year studies, leads to behaviour which is nearly optimal over a longer time horizon. In other words, a good basis is one where pursuit of short-term objectives leads to long-term optimal behaviour. To the extent that there is a conflict between long and short-term objectives, it is the valuation basis which is at fault. This is the philosophy behind the use of embedded values; by maximising embedded value earnings over a short term, management also theoretically maximises shareholder value over the long term. Much of the embedded value literature concentrates on how the calculation should be carried out in order that this principle actually holds in practice. A more mathematical treatment is contained in Appendix A.

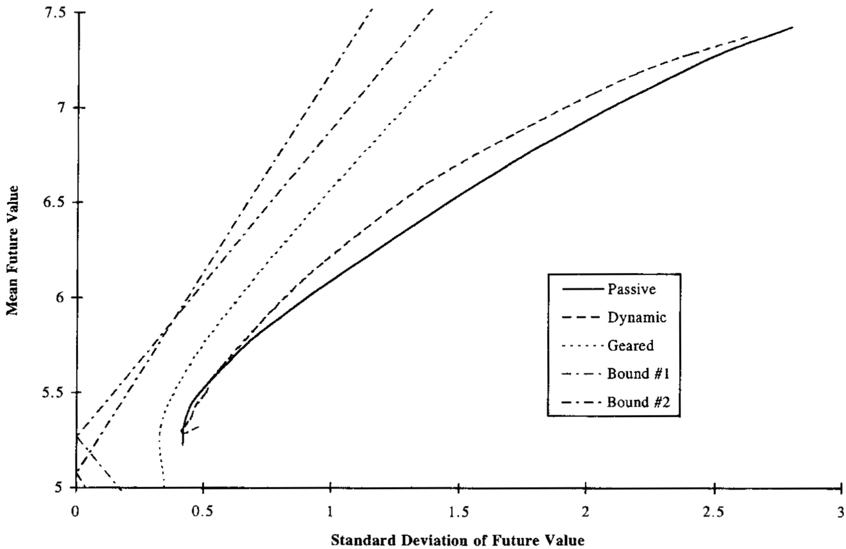

Figure 3.6.1. Passive, dynamic and geared frontiers with dual bounds for the random walk model

3.6 *Results of Dynamic Optimisation*

3.6.1 Turning firstly to the random walk model, we can see, in Figure 3.6.1, that the dynamic optimisation provides only a small improvement in the frontier — indeed, the out-performance is probably not large enough to be economically significant. The more important point which emerges is that, in principle, static strategies may be beaten by dynamic ones. Some of these dynamic strategies take a speculative position when the asset share is less than the reserve, to increase the likelihood of catching up. In other words, the winning strategies lock in speculative gains by selling risky assets if they have performed well. The fact that such a counter-intuitive strategy is optimal in a mean-variance setting is not, I believe, widely appreciated. Other utility definitions would produce the more conventional wisdom that one should increase exposure to risky assets as the surplus rises, and, arguably, such alternative formulations give rise to better decisions. In addition to allowing positive portfolio mixes, I have also calculated the frontier allowing for gearing and short positions, although these are largely discouraged by the 'in connection with' test in the Insurance Companies Regulations 1994. Interestingly, gearing does allow some significant improvement, particularly at the higher-risk end of the risk spectrum. In particular, an investor with a risk appetite equivalent to 100% equity can get a considerably higher expected return by holding a diversified portfolio containing equities, property and gilts, then gearing up by borrowing cash. Somewhat surprisingly, a geared portfolio can also enable the absolute variability to be reduced at the low-risk end of the spectrum. This is largely due to 'silly strategies', including, for example, a strategy which, in the event of exceptionally high investment returns in the first few years, takes short positions on equity markets in order to achieve a compensatingly low return over the later years. I have also plotted some dual upper bounds. We can see that the geared strategies come moderately close to these upper bounds. This suggests that, although our dynamic strategy was not proven optimal in any rigorous sense, there is only limited scope for improvement, even given unlimited resources of ingenuity.

3.6.2 For the chaotic model, we can predict asset returns exactly one year in advance. This enables us to be in the best performing asset class each year. Such a strategy produces mean values around 12, which are off the top of my chart. If we permit short positions, then the return can be made arbitrarily high by shorting underperforming assets to gear up the outperforming ones. Plainly it is nonsense to suggest that the above strategy could be implemented in the real world. However, the existence of such strategies does differentiate chaotic models very strongly from random walk approaches, even if, in a passive framework, the two look rather similar.

3.6.3 For the fractal model, I have not been able to achieve any significant improvements from dynamic trading. Interestingly, the ability to take short positions is also of no use within the fractal model. The charitable strategy eliminated fat upside tails, but nothing can be done about fat downside tails, so that every strategy involving short positions has infinite standard deviations. Arguably the thinner-tailed models, particularly those based on normal distributions, materially understate the importance of the jump exposure which arises from short positions. It is dangerous to assume that all aspects of risk can be captured in a single volatility figure; the result of such an assumption could be the adoption of speculative strategies which go badly wrong in the real world.

3.6.4 The Wilkie model, in Figure 3.6.4, allows substantial improvement from dynamic optimisation, equivalent to around 3% p.a. The main element of the improvement is the detection and exploitation of speculative trading opportunities. In this regard, the Wilkie model falls between the efficient market of the random walk model and the chaotic model where returns are totally predictable. Allowing for short positions, I obtain a further dramatic improvement

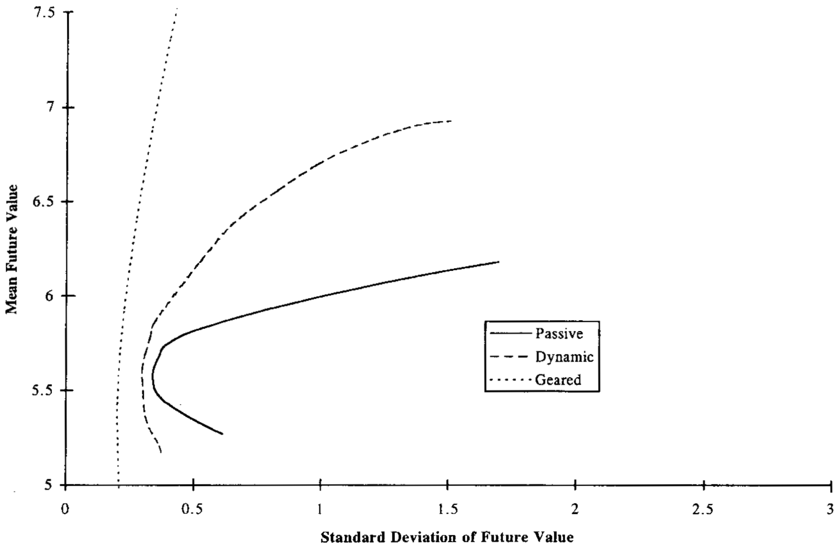

Figure 3.6.4. Passive dynamic and geared frontiers for the Wilkie model

and a significant reduction of the minimum risk position. If this model is correct, then it is very much in the interest of the policyholder to take short positions from time to time, and, if followed wisely, this can dramatically improve performance without increasing the risk. This has significant implications in almost every situation in which such models are applied. The result illustrates an important point. The Wilkie model describes a market which is *inefficient*, that is, assets are mispriced from time to time. However, there are no *arbitrage opportunities*, so that, while some strategies may have abnormal expected returns relative to the risk entailed, none of these opportunities are totally riskless. It could be argued that it is prudent to use a model which captures perceived inefficiencies in the real world. For example, consider an investment manager who buys a stock which seems cheap on fundamentals, but whose market price subsequently collapses. If the fundamental analysis was sound, then this market loss could be attributed to market inefficiency, and due account should be taken of its potential adverse effect on performance. However, Figure 3.6.4 suggests that market inefficiencies are overwhelmingly a good thing for investors, because of the outperformance which is possible from speculative trading. The Wilkie model is particularly helpful in this regard, because the same inefficiencies persist for many years. The use of such a model is, in my view, substantially less prudent than a random walk model, where such exploitable features are absent from the start.

3.6.5 Moving now onto the Dyson & Exley model, shown in Figure 3.6.5, we can see that dynamic optimisation allows significant reduction in risk. This is

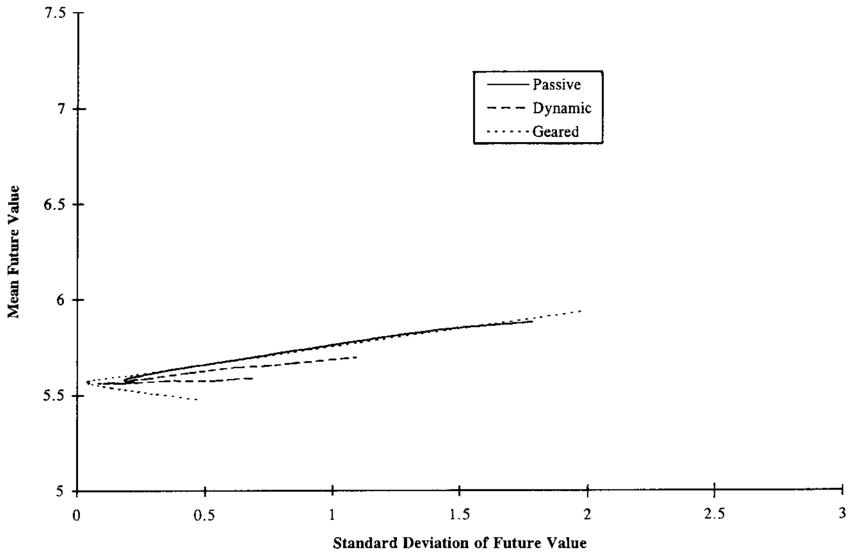

Figure 3.6.5. Passive, dynamic and geared frontiers for the Dyson & Exley model

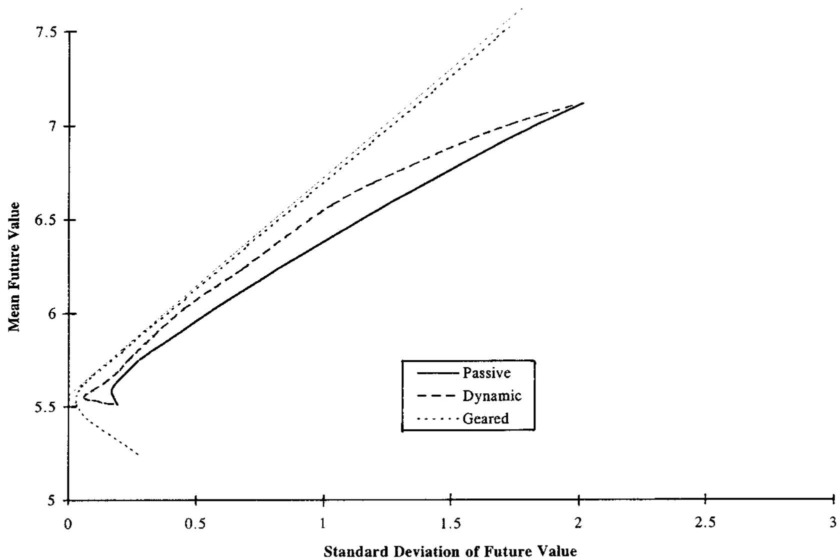

Figure 3.6.6. Passive, dynamic and geared frontiers with dual bounds for the jump equilibrium model, including dual bound

because of an old idea, that of *immunisation*, first described in Redington (1952). At the start of the policy, we have a liability with term 5, and assets which are the future premiums plus whatever investments are held. In order for the mean term of the asset to equal the liabilities, these investments should be held in longer-dated index-linked gilts, moving into shorter-dated stocks as time passes. I did not have to do any of these calculations explicitly; rather, my dynamic strategy works 'blind' without knowing the features of any particular model. The fact that the blind strategy found the immunised position is, therefore, reassuring. It must be regarded as a serious weakness of conventional asset-liability models that strategies such as this, which have been understood for many years, are not revealed as optimal. The reason, of course, is that many of the models do not describe sufficiently many investment vehicles to test the strategies you want to test. Interestingly, at the higher risk end, my dynamic strategy performs very poorly, even underperforming the static approach. This is because the strategy was based on a hunch, and applied in the same manner for each model. In this case my intuition was wrong. The picture, when bonds of several maturities are allowed, is somewhat more alarming. One can obtain arbitrarily high expected returns by borrowing medium term, in order to invest in a mixture of long and short bonds. This so-called 'barbell strategy' is an arbitrage opportunity first described in Redington (1952), and is a simple consequence of the desire to make the convexity of the assets as large as possible and that of the liabilities as small as possible.

3.6.6 Finally, we consider the jump-equilibrium model, shown in Figure 3.6.6. Examining the chart, we can see that dynamic optimisation makes most difference for the low-risk end of the frontier, primarily because of the effectiveness of immunisation. However, in this case the geared strategy comes much closer to the theoretical upper bound. It can be shown that the upper bound is attained by a portfolio involving various derivatives. As one would expect for an equilibrium model, the use of derivatives is actually of rather minor benefit; if this was not the case then the demand for the particular instruments in question would be large, and, in particular, would exceed the net supply of derivatives, which is zero.

4. APPLICATIONS TO A MODEL OF THE BUSINESS CYCLE

4.1 *Reasons for Building the Model*

4.1.1 We have seen how individuals can optimise strategies according to statements of preference relationships, such as those expressed by utility functions. Solution of such problems necessarily entails some non-linear optimisation.

4.1.2 It is important to realise that the management of a company for the benefit of shareholders does not usually entail utility maximisation at the company level. Instead, corporate management should be so focused as to enable shareholders to maximise their utility, allowing, *inter alia,* for the shareholders' ability to diversify some kinds of risk. We will see, in Section 6, that the maximisation of shareholder utility can be interpreted as maximising the value of the company.

4.1.3 We have seen, in Section 2, how a company valuation can be obtained by discounting expected cash flows under an appropriate risk-neutral probability law. The optimisation of company value is, therefore, essentially the maximisation of a linear quantity subject to constraints. There is some room for difference of opinion in the choice of risk-neutral law, just as there is some room for arbitrariness in the choice of discount rates, and indeed, given one it is possible to derive the other. One useful input, when determining the risk-neutral law, is the relationship between net assets and share prices for those companies which are quoted.

4.1.4 In a Modigliani Miller world with no transaction costs, all markets have the same expected return after adjusting for risk, and hence there is no merit from a shareholder value perspective in participating in one market rather than in another. However, if there are entry costs to a particular market, then it is economically reasonable that subsequent expected returns on capital would be anomalously high, given the level of risk; indeed, this should be so in some scenarios in order to justify entry in the first place. The model below can, therefore, be considered as a refinement of the Modigliani Miller result to allow for transaction costs.

4.2 *Model Construction*

4.2.1 I have developed the following model with applications to risk-based capital for non-life insurers firmly in mind. However, it has been pointed out to me that the principles may be appropriate to businesses more generally. An investor's reason for putting capital into a commercial enterprise is to obtain a return on that capital. In most situations, after an initial start-up cost, the return obtained from real production is, in the long run, greater than that which might be obtained on deposit. It is potential return which justifies the establishment of businesses, and by which capitalists will judge one project against another. However, in any market, to increase the supply of a product will tend to force the price down. Thus, a capitalist is faced with the dilemma that, the more capital he invests in the production process, the lower the price of each unit and so the lower the unit profit. This observation is sometimes called the law of diminishing returns. In an insurance context, we interpret the capital employed as the free reserves, where the technical reserves are calculated on a realistic basis and excluding the European Union solvency margin. Stockbrokers often calculate this for composite insurers, starting from net assets as in the annual reports and accounts, and unwinding the implicit margins in reserves by adding any additional embedded value from the life operations. The decision as to how much of the capital should be tied up in statutory prudent margins by writing insurance

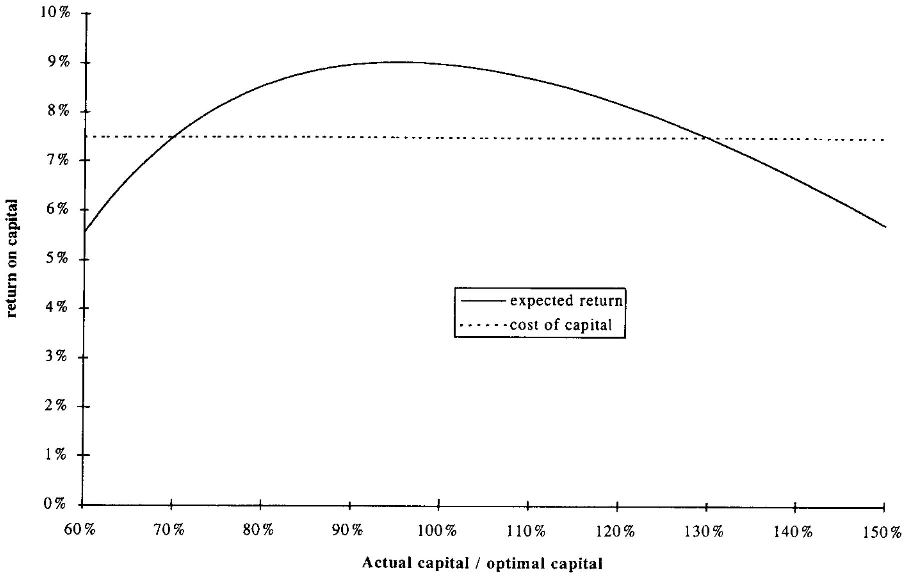

Fig. 4.2.1. Risk - neutral expected return on capital and cost of capital

business, or, alternatively, left in reported free reserves, is an interesting optimisation question in its own right, which we do not consider here. However, we do assume that the result of this optimisation is of the form:

$$\text{Risk neutral expected return on capital} = p\left(\frac{\text{Actual amount of capital held}}{\text{Some measure of market size}}\right)$$

where p is a function describing the profit per unit capital employed. One possible function p is shown in Figure 4.2.1.

4.2.2 There comes a point at which the additional profit available from an increase in production is not sufficient to justify the capital injection required. At this point, the project, in a sense, is using exactly the right amount of capital, and investors will not voluntarily contribute further capital to the enterprise. Conversely, investors would be unwise to take any capital out. We can describe this concept as *capital adequacy* from a shareholder's perspective; plainly this has nothing to do with capital adequacy in the sense of ability to pay claims. The value added from the insurance process per unit time will be the return earned minus the cost of capital, which is:

$$\text{Value added} = u\left[p\left(u/S\right) - r\right]$$

where u is the capital held, S is the market size and r is the cost of capital, which, under the risk-neutral measure, is the risk-free rate. I have taken r to be 7.5% in what follows. The optimal value of u maximises the value added, for which the first order condition is:

$$p(x) + xp'(x) = r$$

where $u = xS$. Notice that, in general, the optimal value of x is slightly higher than the value of x which maximises $p(x)$. There is some arbitrariness in the measure of market size, and we scale the definition of S such that the above equation is solved at x=1. We call S the *optimal level of capital*. Using the above condition, we can define quantities α and β by the Taylor series:

$$p(x) = \alpha + (r - \alpha)(x - 1) + (\alpha - \beta - r)(x - 1)^2 + O(x - 1)^3$$

or, equivalently:

$$xp(x) = \alpha + r(x - 1) - \beta(x - 1)^2 + O(x - 1)^3.$$

We can interpret α as the return on the optimal amount of capital, while β describes the extent to which deviations of capital from the optimal amount lead to decreased returns. The 'break-even' values of x, for which the capital is only just serviced in the short term, are given to first order by $x = 1 \pm \sqrt{\frac{\alpha - r}{\beta}}$. I have taken α = 9% and a break-even range of [70%, 130%], implying β=0.166. The purpose of this proposed model is to construct the optimal amount of capital from a small number of assumptions, and to investigate the means by which it might be maintained at or near this level. The traditional approach to capital allocation is based on the theory of solvency, which is the ability to pay trade creditors as they fall due. The good capitalist has little concern for the welfare of these creditors, following the sole objective of added shareholder value. This concept is not new, and the maximisation of appraised values is suggested, for example, in Mehta (1992) or Bride & Lomax (1994).

4.2.3 Over an economic cycle, the inherent profitability of a particular enterprise may vary. It follows that the optimal amount of capital will also change from time to time. The capitalist making rational investment decisions will then find himself trying to hit a moving target. This uncertainty is compounded, as the profits earned by the enterprise are also highly variable. In theory, the perfect capitalist will continually be either extracting dividends or subscribing to new issues in order to maintain capital at its ideal level. Let us suppose that the market size grows at a (risk-neutral) expected rate μ. If the capital is maintained at its optimal level, then the value added per unit time is αS. Part of this, equal on average to μS, is retained, so that the capital remains at its optimal level, and the remainder, $(\alpha-\mu)S$, is distributed as dividend. The value of the firm is then simply the value of an increasing annuity discounted at the risk-free rate, which is $\left(\frac{\alpha-\mu}{r-\mu}\right)S$. More generally, if the current level of capital u is different from S, then the required cash flow to restore the optimal level of capital must also be included, giving a value of $u+\left(\frac{\alpha-r}{r-\mu}\right)S$. We assume that $\alpha > r$, so that the return on insurance is better than the risk-free rate. This does not contradict the risk neutrality assumption that the expected return on all investments is equal to the risk-free rates, since the insurance business is not a self-contained investment transaction by the shareholders. In order to work back to the risk-free rate, we would need to take into account the entry costs of penetrating the market in the first place.

4.2.4 In my calculations, I have worked on the basis of r =7.5% and μ=2.5%. I have also assumed that, at the optimal capital position, market value is 130% of capital, which implies that α=9%. Then the value of the firm is net assets plus 30% of the optimal capital level. On the basis of historic share prices in the U.K., this would be a rather bullish valuation for an insurer. However, this assumes that capital is always maintained at its optimal level. Deductions must be made from the 130% to allow for sub-optimal capital levels and also for the cost of remedying the situation.

4.3 *Consideration of Transaction Costs*

4.3.1 In reality, it is not easy for a corporation to alternate frequently between capital raising and distribution. The reason is that both of these activities have a cost associated with them. The raising of capital involves investment banks and a raft of other professionals, who extract fees from the proceeds. By the same token, it can be argued that a dividend payment triggers an early payment of corporation tax, which could otherwise have been deferred. An enterprise which makes excessive use of either facility will be a joy to bankers and the tax man, but will be cursed by its shareholders. A more moderate approach is to calculate a comfort zone around the optimal amount of capital. If the shareholders' capital, including retained earnings, is within the comfort band, then no particular action is required to adjust it. In this example, I have taken account only of capital raising costs, which I assume are a proportion λ of the amount of capital raised. I have used λ=3%.

4.3.2 If the level of available capital moves out of the top of the comfort zone, then the time has come to distribute a dividend. This happens at the point where the excess capital is depressing profits to such an extent that it is worth paying advance corporation tax in order to shed some weight. On the other hand, if the available capital moves out of the bottom of the comfort zone, then some serious profit opportunities may be missed. In this situation the shortage of

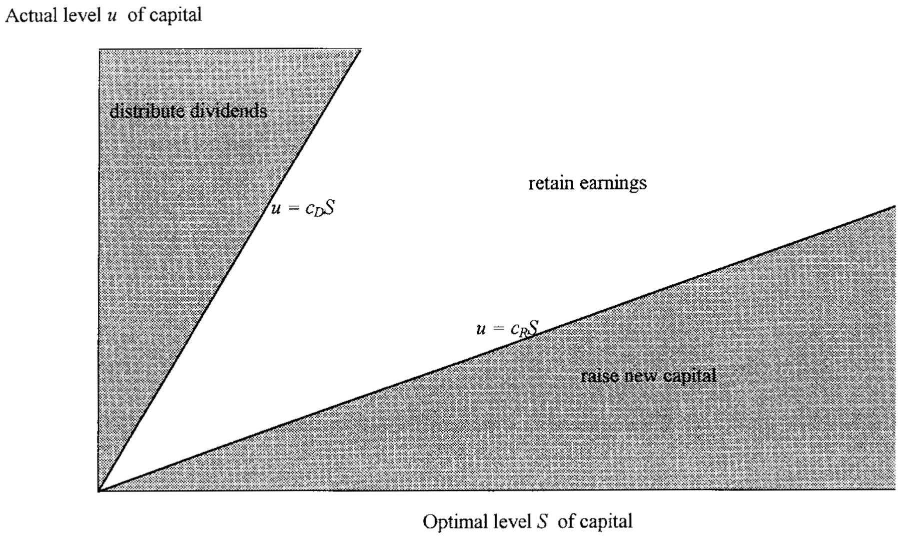

Figure 4.3.2. Three regions for capital strategy

capital is propping up unit profits, and there is good cause to seek more capital to plough back into the business. Indeed, the opportunity is so good as to justify a little extravagance on merchant bankers and other professional advisers. We suppose that dividends are paid when $u > c_D S$ and capital is raised when $u < c_R S$. This divides the possible environment into three regions, as shown in Figure 4.3.2. In each case, the quantum of the capital transaction is just sufficient to get u into the comfort zone $c_R S < u < c_D S$ and no further, so that both dividends and capital raising dribble out in little bits. Of course real life is different, because dividends and capital raising are lumpy. One reason for this is that not all costs are proportional; fixed costs must be taken into account, and this implies that lumps are better than dribbles. However, for many applications the proportional costs are most significant, so we continue with this simplifying assumption to obtain an approximation of how the world should work.

4.3.3 We will need to model the uncertainty of the insurance process. We measure this uncertainty, not in absolute terms, but relative to the ideal amount of capital in the market, so that the insurance process is low risk if the returns obtained allow the company to track the ideal level of capital easily. In the usual situation, where low returns are associated with reduced capacity and hence a rising ideal size, the process may be regarded as more risky. We assume that the volatility of returns relative to the market size is σ per unit time, so that, within the comfort zone where no capital transactions take place, we have the stochastic differential equation:

$$\frac{du}{u} = p\left(\frac{u}{S}\right)dt + \frac{dS}{S} - \mu dt + \sigma dz$$

where z is a Wiener process, or Brownian motion. I take σ as 20% in my base case, and then examine what happens to shareholder value if, by some matching exercise, we can reduce it to 15%. We do not need to assume anything about the noise in the S process or its correlation with z. This model makes the heroic assumption of the absence of jumps in u relative to S. For a given comfort zone, we would expect transaction costs to be more significant if the business is more volatile, for then the situations in which capital is raised will be more frequent.

4.3.4 It now remains to quantify the amount of transaction costs paid in capital raising. Using a local time argument, outlined further in ¶4.3.5, we can see that the first order term for the present value of transaction costs is $\frac{\sigma^2\lambda}{2(r-\mu)(c_D - c_R)}S$. This expression looks reasonable; as anticipated it increases as the volatility σ increases and it is proportional to λ. In addition, it decreases as the width of the comfort zone increases. This again is obviously sensible, as when the band is wider, the boundaries will be hit less frequently.

4.3.5 For the benefit of the mathematically inclined, I shall outline a derivation of the expression in ¶4.3.4 using a local time argument. This is instructive, but also somewhat involved, and is not necessary for understanding the rest of the paper. We first argue that, since the diffusion processes are of infinite variation, the transaction costs would increase without limit as the no-transaction band gets arbitrarily narrow. Now, the difference between the amount of capital raised and dividends paid in unit time corresponds roughly to the excess drift $\alpha - \mu$, and this does not grow as the band gets narrow. Hence, the leading term in the expression for transaction costs will depend on the volatility, while the drift will only contribute to the second order and higher terms. To calculate the first order terms, we can thus take the drift as zero. I now claim that we can approximate x by a function $h(w)$ of a Wiener process w. The function h is the saw-tooth function shown in Figure 4.3.5. Plainly $h(w)$ has the right volatility, namely σ, and also has zero drift between the boundaries. We also note, for later use, that, to first order, this will result in x being uniformly distributed between c_R and c_D. On the boundaries, the drift is exactly enough to make the process reflect, and we now quantify these drift terms, for these are the capital flows we are interested in. Differentiating h twice, we have:

$$h''(w) = 2\sigma \sum_{n=-\infty}^{\infty} (-1)^n \delta\left(w - n\left(\frac{c_D - c_R}{\sigma}\right)\right)$$

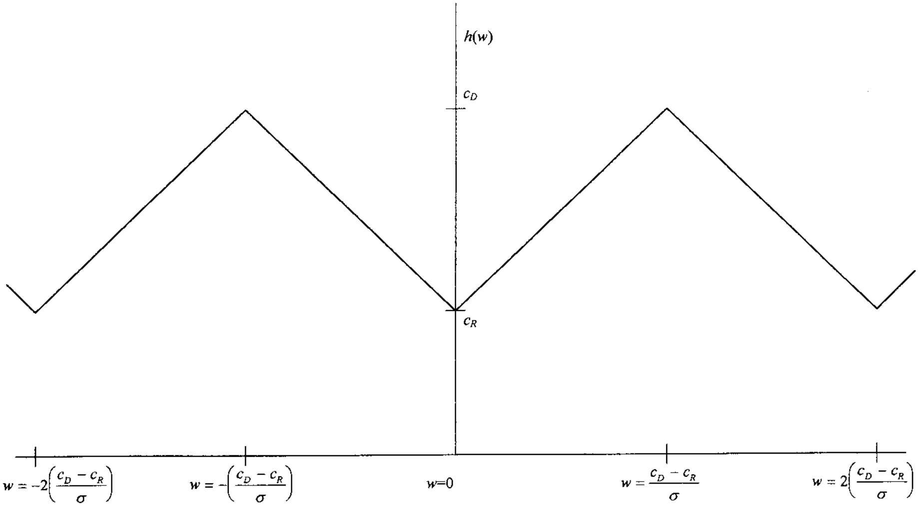

Figure 4.3.5. The saw - tooth function $h(w)$

where δ is the Dirac function. Then, from Itô's formula, we have:

$$x_T = x_0 + \int_0^T \pm\sigma\, dw_t + \sigma \sum_{n \text{ even}} \int_0^T \delta\left(w_t - n\left(\frac{c_D - c_R}{\sigma}\right)\right) dt$$

$$-\sigma \sum_{n \text{ odd}} \int_0^T \delta\left(w_t - n\left(\frac{c_D - c_R}{\sigma}\right)\right) dt.$$

This expresses x neatly as σ multiplied by another Wiener process (which turns out, in fact, to be our original process z) plus the capital injections less the dividends paid. If we look at the second term, which describes the capital injections, we must integrate a δ function each time the Wiener process touches an even multiple of $\frac{c_D - c_R}{\sigma}$, which is, roughly speaking, the amount of time spent at these points, sometimes called *local time*. This concept is described much more rigorously in Rogers & Williams (1987). We can now argue that the proportion of a unit time spent at such points will, on average, be the reciprocal of the point spacing, since all time must be spent somewhere. This suggests the approximation, for large T:

$$\sum_{n \text{ even}} \int_0^T \delta\left(w_t - n\left(\frac{c_D - c_R}{\sigma}\right)\right) dt \approx \frac{\sigma}{2(c_D - c_R)} T.$$

Multiplying by σ gives an approximation for the total capital raised in an interval per unit S. Projecting and discounting give the claimed result.

4.4 *The Effect of the Business Cycle on Company Value*

4.4.1 It is easy to construct a theory of business cycles in terms of a herd instinct, or delayed overreaction to changes in experience. Unfortunately, such a theory usually requires the market participants to ignore emerging information for long periods of time, and having finally assimilated the information into the management process, to do their sums consistently incorrectly; an observation which may reasonably be denied, at least in certain cases. However, cyclical behaviour can also be generated by rational decisions if there are frictional costs in the capital markets. Following a period of healthy profits, the capital level may be near the top of the comfort band. This generates a good dividend flow, but the excessive supply of capital forces a price war. A period of poor profitability ensues, during which much of the capital base is eroded by losses. The weakened capital base restricts supply from the manufacturing base, pushing prices up. At this stage, enterprises will seek to raise more capital to take advantage of the rising rates, and, in doing so, generate good profits, starting the cycle again.

4.4.2 The effect of the cycle is to reduce shareholder value, because the capital is available to support large volumes of new business exactly when the business is least profitable. Conversely, when prices are high a capital shortage is likely to restrict capacity. Thus, the weighted average value added is less than what it would be if exposures were flat over the cycle. We can use this observation to say intelligent things about the value of the insurance company. The starting point is the value $u+\left(\frac{\alpha-r}{r-\mu}\right)S$, determined in ¶4.2.3. We must then reduce this to allow for the effect of the cycle and for transaction costs in raising capital.

4.4.3 We can base a theory of shareholder value on discounted expected cash flow. The cash flows to be included are the dividends minus additional capital subscribed, and the whole exercise is performed under a risk-neutral probability. The value of the enterprise is then a function of how wide the comfort zone is chosen to be. The optimal zone is then the one which maximises the value of the enterprise. Other aspects of the business, such as the territorial mix or the investment portfolio, will also affect the parameters of the model, and hence the shareholder value. This provides a means for measuring the value added from an asset-liability study, as promised in ¶3.1.4.2.

4.4.4 We have already calculated the present value of expected transaction costs. Let us now concentrate on the cost of having non-optimal levels of capital. We have seen, in ¶4.2.2, that the value added from insurance is $xS[p(x)\text{-}r]$, where x is the capital divided by the ideal level of capital. If x is constrained between c_R and c_D, then, on average, the value added will be, expanding to second order:

$$\frac{S}{c_D - c_R}\int_{c_R}^{c_D} x[p(x)-r]dx \approx \left[\alpha - \frac{\beta}{3}\left((c_D-1)^2+(1-c_R)^2-(c_D-1)(1-c_R)\right)\right]S.$$

Growing this at a rate μ and discounting, we obtain the present value of the insurance value added. The first term is what we had before, while the second term captures the cost of the cycle. The second term is:

$$\text{Cost of cycle} = \frac{\beta}{3(r-\mu)}\left[(c_D-1)^2+(1-c_R)^2-(c_D-1)(1-c_R)\right]S.$$

Naturally this cost decreases as c_D and c_R get close to 1, for, in such circumstances, the capital markets are frequently accessed to ensure that the level of capital stays at the optimum.

4.5 *The Benefits of Risk Management*

4.5.1 We can add up the present value of the transaction costs and the cost of the cycle, expressed in units of S, to give:

$$\frac{\sigma^2\lambda}{2(r-\mu)(c_D-c_R)} + \frac{\beta}{3(r-\mu)}\left[(c_D-1)^2+(1-c_R)^2-(c_D-1)(1-c_R)\right].$$

We wish to choose c_R and c_D to minimise this total cost, and hence maximise shareholder value. It is easily confirmed by differentiation that this cost is minimised when:

$$c_D = 1 + \tfrac{1}{2}\sqrt[3]{\frac{3\sigma^2\lambda}{\beta}} \qquad c_R = 1 - \tfrac{1}{2}\sqrt[3]{\frac{3\sigma^2\lambda}{\beta}}.$$

The minimised cost is $\frac{1}{4(r-\mu)}(3\sigma^2\lambda)^{\frac{2}{3}} \times \sqrt[3]{\beta}$. In my base case I find c_D = 1.14 and c_R = 0.86, so that the width of the 'no-transaction' region is 28% of the optimal capital level. Assuming net assets are currently at their optimum, the value of transaction costs is 4.3% of net assets, while the cycle costs 2.1% on a present value basis. Thus, the total goodwill is around 23.5% of net assets. If I manage to reduce σ from 20% to 15% using an asset-liability model, then goodwill rises to 25.6%, so that the matching exercise has added 2% to the shareholder value.

4.5.2 There are various devices which a company can employ in order to alter the risk-return characteristics of profits. For example, it may use derivatives to hedge market risks, or may buy insurance or reinsurance against event risks. It is sometimes taken as self evident that any enterprise will wish to reduce fluctuation in profits for the benefit of its proprietors. On the other hand, one can argue that any market risk in a commercial enterprise is passed straight on to the shareholders. If shareholders are averse to this risk, they may be prepared to pay some fees to hedge it, while, if they are prepared to take the risk, they will not hedge. All that is required is for enterprises to communicate to shareholders the market positions that they are taking. The study of capital enables issues such as this to be placed in a more quantitative framework. There are good reasons for enterprises to maintain their capital near an optimal level, thus maximising the value added by core operations. A risk management policy can reduce the need for capital raising, so cutting costs and adding value to shareholders. It now becomes clear that reducing risk in an absolute sense is only half the story. The purpose of risk management should be to match investment performance to future capital requirements; the value to shareholders of such preventative measures is described by the formula above. Capital requirements, in this sense, are not so much a function of existing business as of new business opportunities.

4.6 *Improving the Model*

4.6.1 Appraised value calculations would usually contain a far more detailed analysis than I have presented here of a single scenario. The additional detail consists mainly of linear terms, so does not affect the essential nature of the problem. This model has captured something which most appraised value calculations do not, that is the costs which arise in scenarios we would rather avoid (unless one argues that they are all somehow subsumed within the risk discount rate). Minimising these costs gives a sound basis for capital management strategies; strategies whose optimality would be far from evident from a traditional appraised value calculation.

4.6.2 When approached rigorously, the problem that I have considered comes down to a movable boundary partial differential equation. Such problems are not, in general, tractable analytically, and are also computationally intensive to integrate numerically. However, asymptotic analysis of such problems is often fruitful, and the formulae I have shown can be derived more rigorously as the leading terms in asymptotic expansion.

4.6.3 The results here generalise considerably, and the mathematics which arises is very similar to the mathematics of option pricing with transaction costs, as in Davis, Panas & Zariphopoulou (1993). A delightfully elegant asymptotic analysis has been carried out by Whalley & Wilmott (1994), who obtain no-transaction bands proportional to the cube root of the transaction costs, as I have done. However, they take the expansion to higher order terms and in far greater generality than I have done here.

4.6.4 The bands of the 'no transaction' region may appear rather wide, at 28% of net assets. The actuary who follows my approach may appear to have let go of the reins. From another perspective, the wide region reflects the fact that market values are volatile, and so it is costly to make frequent adjustments as a result of every move in market value. This is something which actuaries have known for many years, and fair grounds for criticising some asset-liability methods based only on market value. A conventional actuarial approach to avoid excessive intervention would be to smooth asset values, so that action is only taken after a prolonged change in market value. A simpler alternative is to use market value, and to tolerate wider movements before action is taken, which is what I have done here. The latter approach turns out to be optimal in terms of maximising shareholder value.

5. APPLICATIONS TO THE CONSTRUCTION OF ASSET MODELS

5.1 *Objectives of the Model*

5.1.1 We have explored a number of problems using stochastic models. Many of the published models consider bonds as an asset class without distinguishing by term. For my purposes, I wanted a model which provided full term structures of both real and nominal interest rates in an economically consistent fashion. I also describe equities and property total returns.

5.1.2 I was also rather concerned that models based on the normal distribution did not seem to produce enough nasty surprises, such as the October 1987 crash. Most published term structure models are based on the normal distribution, so automatically exclude jumps. I decided to build my own model allowing for a term structure, with occasional price jumps.

5.1.3 In addition to the investment aspects of the model, I wanted to have a reasonable chance of using the model to solve optimisation problems. In particular, this meant having a convenient risk-neutral law. Now it can be seen that changing accounting currency effects a corresponding transformation to the risk-neutral law, if assets are to be priced consistently in the two currencies. Of course the true probabilities of events are not changed by such transformation. The possibility exists, therefore, that, by choosing a suitable (and artificial) accounting currency, one may ensure that the risk-neutral law is actually the same as the true law. This is exactly what I have done. Furthermore, I have adopted a convention that the accounting currency includes an allowance for reinvested income, so that, when cash flows are expressed in this currency, values are obtained simply by taking expectations.

5.1.4 I perceived substantial gains from making the model symmetrical. I used the same algebraic model for each asset class. One reason for insisting on symmetry is that it is much easier to bolt together a multi-asset model from its constituent parts. Symmetry also makes later extensions easier. It would be frustrating to build an economic theory around real returns relative to the RPI in the U.K., only to be forced back to the drawing board when a client asks me to build a model for overseas bonds.

5.1.5 Those who wish to identify discrepancies between my model and the real world will not have to look far to find them. The model that I have ended up with is a compromise between theoretical concerns, faithfulness to the data, analytical tractability, ease of construction and use, development budgets and deadlines. There are many effects which I would have liked to capture, but did not. For example, it is often observed that markets tend to have bursts of high volatility alternating with more stable periods, sometimes referred to as an ARCH (auto-regressive conditional heteroscedastic) effect, which I have ignored. Neither have I allowed for mean reverting or error-correction effects. I have little doubt that the gamma distribution can be shown to be a poor fit to the various series where I have used it. Over a short time period I only allow the yield curves to make parallel shifts, ignoring, for example, possible changes in the slope. The model also fails to satisfy some elementary constraints — for example, nominal interest rates should not, in principle, be allowed to become negative because of the option to hold physical cash, whereas real yields on index-linked stocks could go negative in principle. My compromise was to allow negative yields as a possibility for all asset classes, so that I could use the same algebraic form in each case. Finally, many of the fitting techniques that I have used are of questionable validity. Some of these shortcomings are not unique to my model. If any reader spots further weaknesses in the model, I would like to know about them. It is up to the user to judge whether the drawbacks of a particular model are sufficiently severe that the whole modelling process is futile, and if so, what is to be done instead.

5.2 *Generation of Jump Processes*

5.2.1 There are various ways of allowing for random variations in simulated series. By far the most popular has been the use of normally distributed random variables as error terms, which gives rise in continuous time to what have been called *diffusion processes.* One good reason for making this choice is the existence of a large volume of literature on such processes. In particular, the application of Itô calculus to diffusion models has been hugely influential in financial economics. Duffie (1992) provides an introduction to these methods. However, I have not used diffusion methods for my jump equilibrium model.

5.2.2 An alternative class of models are compound Poisson processes. These are often employed in collective risk theory to model claims in non-life insurance. Claim events are assumed to follow a Poisson process, while the amounts of each claim are drawn from a specified severity distribution, independent of other claims and also of the underlying Poisson process. Mehta (1995) demonstrates some results from such a process; they have not hitherto been widely used for asset modelling, and I think that they deserve a wider airing.

5.2.3 The model that I have used is a limiting case of a compound Poisson process applied to log asset prices, where infinitely small jumps occur with infinite frequency. At the same time, my model also allows finite jumps with a finite frequency, in contrast to diffusion models. We construct the asset model from a series of jump processes $G_j(t)$ for time $t > 0$ and $j = 1,2,3,4$. The distribution of increments in G_j is given by:

$$G_j(u) - G_j(t) \sim \Gamma[\alpha_j(u - t)]; \quad j=1,2,3,4$$

where α_j is a positive parameter and $\Gamma(p)$ denotes the gamma distribution with parameter $p>0$, which has probability density function:

$$f(x) = \frac{x^{p-1}e^{-x}}{\Gamma(p)}.$$

These processes $G_j(t)$ all start at zero, and increments over non-overlapping intervals are statistically independent. The additive property of the Γ distribution allows us to do this. Note that, in terms of collective risk theory, the gamma distribution is taking the place of the aggregate claims distribution, not of the individual claims distribution. These processes G_j are independent for different j.

5.2.4 We also want to allow for random trend terms. One way to obtain random trend terms is by integrating jump processes. This leads us to define:

$$H_j(t) = \int_0^t G_j(\xi)d\xi; \quad j = 1,2,3,4.$$

5.3 *Asset Prices and Term Structures*

5.3.1 We consider the following four asset classes:

P_1 = sterling;
P_2 = the basket of goods forming the U.K. RPI;
P_3 = a basket of U.K. equities, with dividends reinvested; and
P_4 = a portfolio of U.K. property, with rents reinvested.

Bonds are described in connection with the currency in which they are denominated, so that conventional gilts are derived from the dynamics of P_1, while index-linked gilts are based on the dynamics of P_2.

5.3.2 We have adopted an accounting currency, as described in ¶5.1.3, where all assets are valued by taking an expectation. We will determine this currency more precisely later. Since the asset classes are denominated consistently in the same unknown currency, we can eliminate the effect of the unknown exchange rate by taking ratios. Thus, P_2/P_1 is the numerical value of the RPI (in sterling), P_3/P_1 is an index of total equity returns and P_4/P_2 is an index of total property returns expressed in real terms.

5.3.3 The model for each asset class is:

$$P_i(t) = f_i(t)\exp\left[\sum_{j=1}^{4}\left\{\beta_{ij}G_j(t) - \gamma_{ij}H_j(t)\right\}\right]$$

where β_{ij} and γ_{ij} are parameters to be determined, and $f_i(t)$ is a deterministic function of t, also to be determined. For reasons which will later become apparent, we insist that $\beta_{ij}<1$ and $\gamma_{ij} \geq 0$.

5.3.4 We have deliberately chosen our accounting currency so that the true law is also a risk-neutral law. We assume that bonds trade frictionlessly. Taking expectations, we can deduce that the price at time t in units of $P_i(t)$ of a (zero coupon) bond paying $P_i(u)$ at time u is given by the conditional expectation:

$$\frac{E_t[P_i(u)]}{P_i(t)} = \frac{f_i(u)}{f_i(t)} \prod_{j=1}^{4} \exp\left[\left(\alpha_j - \gamma_{ij} G_j(t)\right)(u-t)\right] \left[\frac{(1-\beta_{ij})^{1-\beta_{ij}}}{\left(1-\beta_{ij}+\gamma_{ij}(u-t)\right)^{1-\beta_{ij}+\gamma_{ij}(u-t)}}\right]^{\frac{\alpha_j}{\gamma_{ij}}}.$$

In the special case where $\gamma_{ij} = 0$, the expression inside the product is replaced by $(1-\beta_{ij})^{-\alpha_j(u-t)}$, which is the limiting form. Coupon bonds are priced by adding up the value of each coupon as if they were individual bonds in their own right, plus the value of the principal. The expression above may look like a version of the rational expectations hypothesis. However, we have not assumed any form of market efficiency when coming to this result. The reasons for insisting that $\beta_{ij}<1$ and $\gamma_{ij} \geq 0$ are now clear — they are necessary in order for the above expectations to be finite.

5.3.5 Taking the above expression at $t = 0$, and provided we know α_j, β_{ij} and γ_{ij}, this enables us to determine $f(u)/f(0)$ for each asset class from the known initial term structure. To apply this in the case of equities and property, where we have only measured total return, we can argue that all bonds whose payoff is determined with reference to a total return index would have prices exactly equal to 1 in the absence of arbitrage. The values of $f(0)$ are calibrated to the initial spot prices.

5.3.6 It is a simple exercise to determine the return obtained from holding zero coupon bonds, but constantly rebalancing to retain a constant maturity τ. All the terms in $f(t)$ cancel, and we obtain the total return index:

$$RP_i(t,\tau) = RP_i(0,\tau) \prod_{j=1}^{4} (1-\beta_{ij}+\gamma_{ij}\tau)^{\alpha_j t} \exp\left(\sum_{j=1}^{4} (\beta_{ij}-\gamma_{ij}\tau) G_j(t)\right).$$

5.4 *Properties of the Simulated Series*

5.4.1 We can calculate expected returns analytically. The conditional expectation in ¶5.3.4 applies after a simple substitution to the ratio of two prices. It should be noted that expected returns may become infinite for long time horizons, because of the possibility of continued devaluation of the accounting currency.

5.4.2 It is also of interest to determine the properties of log total return series, and, in particular, the ratios of such series. This gives series of the form:

$$\log\left(\frac{RP_2(t,\tau_2)}{RP_1(t,\tau_1)}\right)=\log\left(\frac{RP_2(0,\tau_2)}{RP_1(0,\tau_1)}\right)+\sum_{j=1}^{4}\alpha_j t\left[\log\left(\frac{1-\beta_{2j}+\gamma_{2j}\tau_2}{1-\beta_{1j}+\gamma_{1j}\tau_1}\right)\right]$$

$$+\sum_{j=1}^{4}(\beta_{1j}-\beta_{2j}-\gamma_{1j}\tau_1+\gamma_{2j}\tau_2)G_j(t).$$

We notice that each of these series has changes which are independent and identically distributed. We notice also that second and higher moments do not depend on the absolute values of the β_{ij}, but on differences between them (for fixed j and varying i). These observations are extremely useful, because they enable us to separate the calibration of the noise terms from the drift terms. In essence, I have derived the noise terms using statistical techniques, while I have used more economics to get the drifts. The calibration is discussed further in Appendix B.

5.5 *Differences of Opinion*

5.5.1 Many models of the market make the assumption that all investors have the same view of the world, and, in particular, of the distribution of returns on various asset classes. This assertion could be disputed, but it is certainly the case that all participants in the market see the same market prices, irrespective of their view of future market movements. It is, therefore, of interest to consider how differences in market views could affect market prices in the economic model we have constructed.

5.5.2 One interesting question is to consider how differing views of market participants would be reflected in market prices. For example, we may consider two investors whose views differ regarding the likely size of jumps. More specifically, let us suppose that investor A believes the model we have just constructed. One the other hand, investor B does not believe that each of A's processes G_j is really a gamma process as we have described it. Instead, B believes that A's process G_j is actually a gamma process multiplied by a constant factor (which may vary for different j). In effect, B believes in a model of the same form as A, but with scaled values of β_{ij} and γ_{ij}. Working through the algebra, we see that A and B will still expect to see exactly the same prices for everything, after allowing for the fact that their accounting currencies are different. This means that a model such as the jump equilibrium model may arise, even when market participants do not agree on a common 'true' economic model.

5.5.3 Another corollary is that it is not possible, even in principle, to deduce a true probability law from observed market prices. In the example of ¶5.5.2, players A and B believed different probability laws, both of which were internally consistent, but saw the same market prices.

6. THE EFFECT OF TRANSACTION COSTS

6.1 *Model Extensions and Transaction Costs*

6.1.1 Wilkie (1985) considers ways in which his model may be adapted to allow for transaction costs. His suggestion is that for shares, the ask value should be 3% above the 'Wilkie value' produced by his model, while the bid should be 2% below. The bid-ask spread for consols is 1% either side of the Wilkie value. One consequence of these costs is to substantially reduce the returns from trading strategies involving frequent switching. The approach of setting fixed spreads about a mid-market value is widely used by practitioners in other financial contexts.

6.1.2 Geoghegan *et al.* (1992) consider properties which might be expected of call option prices calculated using the Wilkie model, compared to the more conventional Black-Scholes approach. They suggest that the Wilkie model will produce lower call prices (than Black-Scholes) when the equity market looks dear, and higher call prices when the market looks cheap. This would be the consequence of a simple discounted cash flow valuation. One difficulty of this approach is that, applied to the underlying equity, a value is obtained which differs from the market value. The authors could, alternatively, have proposed an option pricing formula which was consistent with today's equity prices, but inconsistent with current interest rates. We must accept that, when a model describes inefficient markets, there will be identifiable anomalies, or 'fault lines'. If the model is extended to more asset classes, one must make a decision as to which side of the fault line the new class will lie.

6.1.3 The same question can be asked of each of the models discussed in this paper: how can we price derivatives given the stochastic behaviour of the underlying assets? On a more elementary level, one might ask how to derive a term structure of interest rates which is consistent with a stochastic model which does not explicitly cater for it. This is a very common practical issue when applying the Wilkie model in real life. If we consider the underlying stochastic models to be a description of a liquid market, we are effectively trying to ascribe 'market-based' values to assets whose prices are not quoted in the underlying market.

6.1.4 It may not be at all clear, on first reading, why these examples are connected. The reason is that absence of price information can be viewed as a limiting case of large transaction costs. When transaction costs are zero, then the market price provides an unambiguous value. With finite transaction costs, it is conventional to work with mid-market prices, but nothing may actually trade there. Indeed, any price which lies within the market bid-ask spread could, in some circumstances, be justified as the underlying 'market-based' value. Further economic theory is required to justify any chosen point within this range. Sometimes, for example in a crash scenario, the bids are repeatedly hit, while no asks are ever lifted. In this case the ask price is far above the market value actually trading, so that the mid-market value does not represent the market value at all. In extreme cases the ask may disappear, effectively becoming infinite. Then we cannot calculate mid-market prices; all we can say is that the market value is at least as big as the bid. An even more extreme case is where the instrument does not trade at all (or, at least, the price is not described by the model under consideration). Effectively, the ask price is infinitely high, the bid price is infinitely low. The market then provides no direct guidance as to value, and one has to rely entirely on economic assumptions. This is what has happened in ¶¶6.1.2 or 6.1.3, and is, in essence, a limiting case of a wide bid-offer spread. If there are no transaction costs at all, then every possible cash flow is traded, so there is no difficulty with computing values — everything can be read from the market. To the extent that we need to determine values of quantities which are not actively traded, we need to understand markets with transaction costs.

6.1.5 We have approached pricing in terms of risk-neutral laws, as described in Section 2. These abstract probability laws give us a means for calculating prices which lie within the bid-ask spread of quoted assets. It might seem reasonable to apply the same formula to price non-traded cash flows. Unfortunately, without picking a particular risk-neutral law, we cannot take this very much further, because different risk-neutral laws would give different values. There is a rather obvious principle that, if an entity can be decomposed as a basket of traded pieces, then the effective ask price of the entity should not be more than the sum of the ask prices of the bits, since otherwise any participant would synthesise the entity out of its constituent bits in preference to lifting the observed ask price. By the same token, the bid price of the entity should not be less than the total bid prices of the bits. This provides possible ranges for bid and ask prices. In most cases, for any proposed value within this range, we can find a risk-neutral law which will suggest that value. In other words, the existence of a risk-neutral law tells us no new information; in order to obtain tighter bounds on market value, we have to find a way of picking one particular law out of all the possible ones.

6.1.6 It is worth emphasising the extraordinary combination of fortuitous circumstances which leads to uniqueness of option prices in the Black-Scholes world. In such a world, it is possible to synthesise a portfolio out of cash and shares such that the value at some future date behaves very much like that of an option. This portfolio involves well-timed purchase and sale of shares, depending on how the underlying price has moved, but, crucially, does not require any foresight into how the price will move in the future. In the continuous time limit, such *hedging portfolios* can replicate options perfectly. It follows from the 'decomposition into bits' argument that the value of the option must equate to the value of the tracking portfolio at prior times. This argument, essentially an application of Section 2.5, is sometimes called *absence of arbitrage*. If the option were priced differently from the hedging portfolio, then riskless profits could be obtained by simultaneously selling the dearer and buying the cheaper — a so called *arbitrage opportunity*. Arbitrage opportunities can be eliminated by the action of a single optimising individual who exploits them.

6.1.7 The arbitrage argument is reinforced by the risk-neutral approach; there is a unique risk-neutral law in the Black-Scholes world, and so there is a unique option valuation framework which is consistent with observed share prices. The option price does not require any assumptions about the nature of the investors participating in the market, and so is said to be *preference independent*. It is important to emphasise that the ability to generate preference-free prices is the exception rather than the rule. The assumptions underlying the Black-Scholes model are strong, and include the absence of taxes or other transaction costs, continuous trading, no jumps in the share price and constant volatility. If these assumptions are relaxed, then option prices are not uniquely determined by the dynamics of the underlying share price. Equivalently, there is then a range of risk-neutral laws, each of which can give a plausible option pricing framework which is consistent with the rest of the market.

6.1.8 The seminal paper on option pricing with transaction costs is Davis, Panas & Zariphopoulou (1993). They write down equations which determine the price at which an investor may wish to buy or sell options based on assumed investor utility functions. The result depends on the utility functions assumed, and so is preference dependent. The general case of multiple asset classes is extremely complex, and I am not aware of any efficient numerical algorithms in practical use.

6.2 *Principles of Valuation*

6.2.1 The results of Davis, Panas & Zariphopoulou illustrate a general point, that when there is friction in a market, the value of a derivative may not be uniquely determined from the probability law of the underlying assets. Instead, we can only make statements about the marginal value of a derivative to a particular class of investor. The marginal value is the equivalent cash flow today which the investor would exchange for a small quantity of the derivative, while leaving the utility unchanged. This definition gives rise to the awkward possibility that the marginal value of a cash flow stream may be different for different investors. Concepts such as *shareholder value* must be more tightly defined if they are to be useful.

6.2.2 An important economic concept is that of equilibrium. Probably the best known, and most widely used, equilibrium model is the capital asset pricing model. Equilibrium describes a situation where all players are solving optimisation problems, and prices are determined by the actions of these players. This hypothesis of total rationality by all market players is often portrayed as very heroic indeed. Equilibrium models are not useful for answering the question "where are the exploitable opportunities that everybody else has missed?". The assumption that the current situation is the outcome of rational competitor action means that the effect of irrational behaviour is concealed. However, such models are widely used for financial management, and appear to lead to sensible decisions. The use of equilibrium models is consistent with asking "what is it about the nature of the market which means that my firm can make money from this activity when others cannot?". At first, I found it astonishing that such models are any use at all. The only explanation that I can offer is that, in the real world, the combination of numerous arbitrary and irrational decisions at a micro level have enough of a common rational thread that the irrational patterns cancel out and we are left with a world which is consistent on a macro scale with rational behaviour.

6.2.3 It is my belief that, over the next few years, equilibrium models will turn out to be far more useful for general management than for asset selection. The major strength of equilibrium models is the ability to investigate and exploit the effect of structural features in the market which differentiate one player from another. The wholesale asset markets are increasingly transparent, and the advent of derivatives has removed many of the structural barriers to gaining economic exposures. Successive governments have encouraged this trend, in the belief that markets should be liberated for the general good. If there are profit opportunities in such a market, then these opportunities will be available to all players in more or less equal measure. To the extent that opportunities remain, many players are failing to optimise, and such behaviour will not be described by an equilibrium model. For that reason, the reception given to financial economics by practical asset managers has been rather less than enthusiastic. For one example, I quote Clarkson & Plymen (1988), who "are satisfied that [the CAPM] using historic betas for return and risk has no practical application to portfolio management, either to improve performance or to reduce the risk". For general management purposes, the situation is very different. There are many structural features which distinguish players and may impact their ability to generate future profits. Examples are differing brands, distribution channels, data quality, capital structures and tax positions. It is sensible for managers to wish to exploit any of these features when this leads to competitive advantage, and while the optimal behaviour of competitors may limit such exploitation, it does not preclude it altogether. I have constructed an equilibrium model of the personal lines insurance market, which I have now used on a number of assignments. The insight obtained from explicitly modelling competitor actions is far greater than one obtains from a more passive approach such as traditional appraised values. However, my model only admits a market cycle to the extent that it can be explained in terms of rational actions by all individual market participants. If I had intended to measure the exploitability of competitor irrationality, an equilibrium model would be of little use.

6.2.4 There are some situations where an actuarial model needs to incorporate market inefficiencies, for example when making stock selection decisions. One way of constructing such models is to build an equilibrium model and then adjust the model for known or perceived inefficiencies at the date of the investigation. I would usually assume that these inefficiencies are priced out of the market within a time horizon of a few years, and the market behaves efficiently thereafter. An alternative approach, which has been widely recommended, is to fit simple time series models to historic data, so that any past inefficiencies are automatically projected into the future. I have never succeeded in getting useful models this way. What tends to happen is that assets which have outperformed in the past are projected to continue to do so in the future, so that, using such models, institutions are encouraged to buy at the top of the market. While the market may be inefficient in the future, it is not obvious to me why the future inefficiencies should be the same as those which occurred in the past. Mehta (1995) has pointed out "a natural tendency for anomalies to be exploited, and hence disappear, once discovered". Future trading opportunities are more likely to arise from factors which nobody has yet thought of. By merely extrapolating past speculative successes, the actuary overstates the likelihood that his particular fund manager will be able to exploit future market imperfections. An additional, and quite separate, problem is that many of the time series models in the statistical literature are inconsistent with equilibrium for *any* choice of parameters. It is easy to greet modelled inefficiencies with misplaced optimism when the 'detection' of such anomalies was inevitable, whether or not the original data were sampled from an efficient market.

6.2.5 The condition for dynamic optimality, as developed in Appendix A, is that the marginal value of a traded cash flow stream lies between the market bid and ask prices. If there are no transaction costs, then marginal values for all assets will be equal to the single market value. In particular, the marginal value of a particular asset will be the same to all market participants. This is extremely convenient, because it means that we can talk about shareholder value and all agree on what it means. Practitioners calculating shareholder values will often do so *as if* there were no transaction costs. The alternative, that costs are taken into account, leads to a model which is very much more complex, and which would not be able to deliver a single value as output. A client who asks an actuary to compute the shareholder value of an enterprise is implicitly assuming that the calculations will be carried out within a theoretical framework where shareholder value is a meaningful concept, which means no transaction costs. If this assumption is good enough for overall company management, it ought also to be good enough for the actuary producing figures for input to the management process. Of course the justification of conventional discounted cash flow techniques for shareholder value analysis depends on frictionless markets just as much as some of the more exotic models do.

6.2.6 We return briefly to the distinction between static and dynamic optimisation that we first met in Section 3. For a given player and stochastic model, calculating a dynamic optimum is technically more demanding than calculating a static one. When constructing equilibrium models, economists usually assume dynamic optimisation. Why is this — surely it would be simpler to use static optimisation instead? In fact, building equilibrium models from static optimisation gives very unsatisfactory results, because demand is price-inelastic. If there is no mechanism for investors to bias their portfolio towards cheap stocks, there is no buying pressure to stop cheap stocks from becoming even cheaper. Fantastic trading opportunities very soon arise, and, by hypothesis, all investors who are following passive strategies will fail to exploit these opportunities.

6.2.7 In the context of our PEP example in Section 2, there is not a freely traded market in lapse rates or in management expenses. Consequently, it is not possible to calibrate the risk-neutral probability law directly from market prices. Instead, it appears that we must make some assumptions about investor utility for the type of investor who is likely to be a shareholder in the PEP provider. There is strong practitioner resistance to the use of utility functions, and fortunately, in many cases, we can avoid having to use them. I will now attempt to explain why we have recourse to this unexpected luxury.

6.3 *Conditional Expectation Algorithms*

6.3.1 There is a convenient half-way house between efficient markets and inefficient markets, which enables us to derive tractable approaches to valuation problems. The idea is that there is a core financial market within which there are no transaction costs, and price quotes can be obtained for every derivative in these core markets. These markets are not necessarily efficient, although they may be. It is not necessary to know the probability laws driving the core markets, but only the prices of various derivatives, and hence a risk-neutral probability law.

6.3.2 All other cash flows outside the core markets are non-traded. We wish to determine the shareholder values of such non-traded cash flows, or more precisely, the marginal values of these cash flows to a particular shareholder. We assume, further, that the reason that these cash flows do not trade is that nobody really wants to, so that, for example, no mechanism which unbundled the cash flows of an insurer into consistent parts would actually result in a preferred stream of cash flows for each counterparty. If this were not the case, then the unbundling would have already happened due to market demand, and the cash flows would then be part of the core market rather than the non-traded market.

6.3.3 I now consider what would happen if the preference relation is a function only of core variables, and the optimal aggregate portfolio invests only in the core market. One consequence of this assumption is that all shareholders have consistent marginal values, even though their utility functions may be different and there are transaction costs in the market. It, therefore, makes sense to talk about *shareholder value*, meaning a common assessment of marginal value by all optimising shareholders.

6.3.4 It then follows that the marginal value of a cash flow stream is the same as the marginal value of a transformed stream, where the transformation is the conditional expectation of the cash flows given the core variables. However, by definition, the conditional expectation is a function only of the core variables, and so the transformed variable is a derivative on the core markets whose value can be observed in the market place. This provides a total consistent methodology for pricing all cash flows.

6.3.5 We now consider two worked examples, based on the PEP appraised value from Section 2. We consider first the lapse risk, and secondly the inflation risk.

6.3.5.1 Working first on the lapse risk, let us suppose that the lapse rate is a decreasing function of investment return, so that, for example, with a 0% return we have 12% expected lapses, but with a 15% return we have only 10% expected lapses. We can interpolate between these values, giving a relationship of the form:

$$\text{Expected lapse rate} = 1 - \alpha * (1 + \text{investment return})^{\beta}$$

with $\alpha = 0.8800$ and $\beta = 0.1608$. If investment markets do well, the office benefits, not only from higher management fees per unit, but also from a larger number of units remaining in force. Instead of the fees being proportional to a total return index S, it is proportional to a $S^{1+\beta}$. We can value these cash flows using an option pricing approach, as in Section 2.6. Based on a volatility of 16% p.a., this suggests that the value of next year's release of reserves on units with current funded value 1 is the management charge multiplied by 0.8926. Subtracting this from unity, we can deduce the risk neutral lapse rate shown in ¶2.4.3.

6.3.5.2 Let us now consider expense valuation as our second example. For the office concerned, it has been observed that the rate of expense inflation per policy has approximately followed RPI inflation, and is unaffected by movements in the equity market. In other words, the expected expenses, given the core variables, is equal to the current expenses rolled up at the RPI. However, the RPI is itself correlated with the equity market. This means that, when the equity market has done well, there is a larger number of units in force, each of which also is likely to generate higher expenses. Although there would be a likely unit expense saving if investment markets performed badly, this saving would apply to a smaller number of units, and so the impact would not be as substantial as might appear at first sight. I have applied a crude option pricing approach based on an annual RPI volatility of 3%. I have also assumed a 'unit gain' model, as described by Wilkie (1995), so that the equity return can be expressed as the rise in the RPI plus other terms uncorrelated with the RPI. This implies that the shareholder value of expenses one year hence on units with current funded value 1 is the current expense level multiplied by 0.8588. This then implies the risk-neutral expense inflation rate shown in ¶2.4.3.

6.4 *Linearisation Algorithms*

6.4.1 The methodology in Section 6.3 is particularly simple to execute if the conditional expectations happen to be linear in the core assets. In this case one is left with the valuation of core assets only, which comes down to knowing the current market value of these core assets. In such cases it is not necessary to obtain derivative prices, which sometimes only exist in theoretical models and not in the real world.

6.4.2 Even when the conditional expectations are not of a linear form, one may still be able to approximate them with such a form. Determining the 'best' approximation is essentially an exercise in linear regression. The closer the approximation, the less the exact form of regression matters. When applied to the asset side, this approach essentially produces the capital asset pricing model, and the multiple regression analogue produces arbitrage pricing theory (APT). The same techniques have also been applied on the liability side; for example Hindley & Smith (1991) analyse financial reinsurance contracts using regression methods.

6.4.3 When approximating a cash flow with its conditional expectation, the residual error has zero conditional mean. This enables cash flows to be decomposed into a 'systematic part', which is a function of the core variables, plus 'specific risk', which has zero mean, given the core variables. It is often argued that specific risk may effectively be diversified at the portfolio level, and should not, therefore, affect pricing, and, indeed, this is a consequence of the assumptions we have made above. By contrast, the regression approach merely produces a residual which is uncorrelated with the core variables, and may still be functionally dependent on them, as occurs, for example, in many option trading strategies. It is by no means obvious that such residuals should not affect market price, although such non-interference would be a consequence of the CAPM. Typically, the numerical differences between the CAPM and more complex pricing approaches are very small.

6.4.4 The linearisation approach does have some inconvenient aspects. One of these is that the regression line changes if the accounting currency changes, because each point is reweighted by the appropriate exchange rate. As a consequence, the value of a series of cash flows may be viewed differently by analysts working in different currencies. The discrepancies are usually small, so that good management is not jeopardised by the use of such approximations. The main inconvenience is that calculations fail the usual accounting reconciliations by a small degree. It is easy to waste time looking for a mistake when the real problem is a mildly inconsistent valuation framework. Similar problems arise when selecting the frequency of revaluations; the product of two returns over successive periods, both linear in the core variables, gives a non-linear return. This, again, can lead to small reconciliation errors which are irritating, but not generally serious. Finally, the use of linear approximations means that the expected market return on core assets may have a small effect on the calculated value, while in the exact model the effect of this assumption is completely self-cancelling.

6.4.5 Financial economics is not the only area where conditional expectations are commonly approximated by linear forms. A familiar application for actuaries is credibility theory. Within certain model classes, a Bayesian posterior risk cost forecast is a weighted average of the prior mean and the experienced risk cost. In other cases, including many encountered in practice, the Bayesian analysis is intractable. In an ingenious paper, Bühlmann & Straub (1970) developed a linear approximation to the conditional mean which works more generally. The subject has become known as 'empirical Bayes credibility theory', and a great deal has been written about it subsequently. In effect, Bühlmann & Straub have applied the APT approach to the pricing of liabilities. The least squares approach to statistical estimates is parallel to the quadratic utility optimisation in the CAPM, and this is why the models are mathematically very similar.

7. Some Concluding Remarks

7.1 Actuaries pre-date financial economics by at least a century, but the problems addressed by the two schools intersect to a large degree. Financial economics has opened a Pandora's box of new concepts and techniques. Many of these techniques were developed in the context of asset selection models. For actuaries, the insight which can be obtained into the behaviour of liabilities and their interaction with assets is at least as important as the asset models themselves. Financial economics is eminently applicable in the traditional fields of actuarial endeavour, namely: insurance, pensions and investments.

7.2 The first models built by financial economists have seen some criticism for being unrealistic, analogous to the weightless strings and frictionless pulleys in high school mechanics. However, economists have not been idle, and many of these restrictions have been relaxed in recent years. As these models become more realistic, the case against using them becomes weaker. However, we should not misunderstand the purpose of economic assumptions. All models are wrong in some respect, but in the real world we have to make decisions. Either we use a flawed model to help, or we use no model and rely on 'feel'. The possibility of using a perfect model does not exist, because nobody has built such a model. It is remarkable how even models which seem blatantly unrealistic can sometimes give rise to good management decisions.

7.3 Many financial models can enable a better understanding of the world about us. The application of models only shows financial rewards when the model is used as an aid to decision making. To make the best decisions, some form of optimisation is involved. Some of the most successful asset models have described asset prices as the outcome of many investors each solving optimisation problems. Optimisation is essential to the role of the actuary, but the necessary techniques have been principally developed by economists.

7.4 Some have commented that this paper is too mathematical for its intended audience. To most actuaries, investment mathematics means compound interest, while financial economists have embraced more advanced techniques, which have turned out to be highly fruitful. It is a matter for regret that actuaries have thus been overtaken, particularly since many actuaries were once fine mathematicians. I hope that this paper will expedite some catching up.

7.5 Sometimes financial economics reinforces the results of more conventional methods. In other applications, financial economics addresses questions which are difficult to express and even harder to solve within the traditional framework. I do not see classical methods as an alternative to new techniques; rather, the alternatives to the use of financial economics are to leave the questions unanswered, or to build a new solution from scratch.

ACKNOWLEDGEMENTS

I am grateful for the considerable help and advice I have received from peers and colleagues in preparing this paper. Special thanks are due to James Bridgwater, Francisco Comprido, Jon Exley, Malcolm Kemp, Patrick Lee, Shyam Mehta, Steve Nagle, Simon Pollack, Steve Satchell and Tim Sheldon for many fruitful discussions. Any errors and omissions are, of course, my own responsibility. Any views expressed are mine alone, and not necessarily those of my employer or anyone else I have fraternised with.

REFERENCES

BARCLAYS DE ZOETE WEDD (1995). *Equity gilts study.*

BELLMAN, R. (1957). *Dynamic programming.* Princeton University Press.

BENJAMIN, S. (1959). The theory of games and its application to interest rates. *J.I.A.* **85**, 373-413.

BENJAMIN, S. (1984). An actuarial layman looks at control theory. *Transactions of the 22nd International Congress of Actuaries*, **4**, 295-309.

BLACK, F. & SCHOLES, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, **81**, 637-654.

BLACK, F., DERMAN, E. & TOY, W. (1990). A one factor model of interest rates and its application to treasury bond options. *Financial Analysts Journal*, 33-39.

BLUNDELL, G. (1995). Investment in property: a victim of the quants? *The Actuary*, **6**, No 2 (October), 12-13.

BOOTH, P.M. (1995). The managment of investment risk for defined contribution pension schemes. *Transactions of the 25th International Congress of Actuaries*, **3**, 105-122.

BRIDE, M. & LOMAX, M.W. (1994). Valuation and corporate management in a non-life insurance company. *J.I.A.* **121**, 363-440.

BÜHLMANN, H. & STRAUB, E. (1970). Glaubwürdigkeit für Schadensätze. *Mitteilungen der Vereinigung Schweizerischer Versicherungs-Mathematiker*, **70**, 111-133.

BURROWS, R.P. & WHITEHEAD, G.H. (1987). The determination of life office appraisal values. *J.I.A.* **114**, 411 - 465.

CHAMBERS, J.M., MALLOWS, C.L. & STUCK, B.W. (1976). A method for simulating stable random variables. *Journal of the American Statistical Association,* **71** (June), 340-344.

CLARK, G. (1992). *Asset and liability modelling - the way ahead?* Staple Inn Actuarial Society.

CLARKSON, R.S. (1995). Asset/liability modelling and the downside approach to risk. *Transactions of the 25th International Congress of Actuaries*, **3**, 139-162.

CLARKSON, R.S. & PLYMEN, J. (1988). Improving the performance of equity portfolios. *J.I.A.* **115**, 631-691.

DAVIS, M.H.A., PANAS, V.G. & ZARIPHOPOULOU, T. (1993). European option pricing with transaction costs. *SIAM Journal of Control and Optimisation*, **31**, 470-493.

DAYKIN, C.D. & HEY, G.B. (1990). *Managing uncertainty in a general insurance company. J.I.A.* **117**, 173-277.

DAYKIN, C.D., PENTIKÄINEN, T. & PESONEN, M. (1994). *Practical risk theory for actuaries.* Chapman & Hall, London.

DUFFIE, D. (1992). *Dynamic asset pricing theory.* Princeton University Press.

DYSON, A.C.L. & EXLEY, C.J. (1995). Pension fund asset valuation and investment. *B.A.J.* **1**, 471-540.

FLANNERY, B.P., PRESS, W.H., TEUKOLSKY, S.A. & VETTERLING, W.T. (1989). *Numerical recipes.* Cambridge University Press.

GARMAN, M. & KOHLHAGEN, S. (1993). Foreign currency option values. *Journal of International Money and Finance*, **2** (December), 231-237

GEOGHEGAN, T.J., CLARKSON, R.S., FELDMAN, K.S., GREEN, S.J., KITTS, A., LAVECKY, J.P., ROSS, F.J.M., SMITH, W.J. & TOUTOUNCHI, A. (1992). Report on the Wilkie stochastic investment model. *J.I.A.* **119**, 173-228.

HARRISON, J.M. & KREPS, D. (1979). Martingales and arbitrage in multi-period securities markets. *Journal of Economic Theory*, **20**, 381-408.

HINDLEY, D.J. & SMITH, A.D. (1993). Quantitative analysis of financial reinsurance. *Proceedings of the 24th ASTIN Colloquium*, Cambridge.

HO, T.S. & LEE, S. (1986). Term structure movements and pricing interest rate contingent claims. *Journal of Finance*, **41**, 1011-29.

HOOKER, N.D., BULMER, J.R., COOPER, S.M., GREEN, P.A. & HINTON, P.H. (1996). Risk-based capital in general insurance. *B.A.J.* **2**, forthcoming.

HULL, J. & WHITE, A. (1990). Pricing interest rate derivative securities. *Review of Financial Studies*, **3**, 573-92.

KEMP, M.H.D. (1997). Actuaries and derivatives. *B.A.J.* **3**, *forthcoming*.

LEE, P.J. (1993). Portfolio selection in the presence of options and the distribution of return of portfolios containing options. *Proceedings of the 3rd AFIR Colloquium* (Rome), **2**, 691 - 708.

LOCKYER, P.R. (1990). *Further applications of stochastic investment models.* Staple Inn Actuarial Society.

MARKOWITZ, H.M. (1987). *Mean-variance analysis in portfolio choice and capital markets.* Blackwell, Oxford.

MEHTA, S.J.B. (1992). Allowing for asset, liability and business risk in the valuation of a life office. *J.I.A.* **119**, 385-455.

MEHTA, S.J.B. (1995). Stochastic investment models. *The Actuary*, **6**, No 4 (December), 14-16.

MEHRA, R. & PRESCOTT, E. (1985). The equity risk premium: a puzzle. *Journal of Monetary Economics*, **18**, 145-161.

MERTON, R.C. (1971). Optimal consumption and portfolio rules in a continuous time model. *Journal of Economic Theory*, **3**, 373-413.

MODIGLIANI, F. & MILLER, M.H. (1958). The cost of capital, corporation finance and the theory of investment. *American Economic Review*, **48**: 261-97.

NOWELL, P.J., MCCUTCHEON, J.J., TOMPKINS, P.D.G., MIDDLEMISS, K.A., SCOTT, P.G., HITCHCOX, A.N., GOFORD, J., DASH, A.C., BATTAMS, E.A., MILNE, R.J.H., GOODWIN, E.M. & CAMPBELL, G. B.L. (1996). The future of the profession. *B.A.J.* **2**, 325–427.

REDINGTON, F.M. (1952). Review of the principles of life office valuation. *J.I.A.* **78**, 286-315.

ROGERS, L.C.G. & WILLIAMS, D. (1987). *Diffusions, Markov processes and martingales,* Vol. 2. Wiley.

ROSS, S.A. (1976). The arbitrage theory of capital asset pricing. *Journal of Economic Theory*, 341.

SCOTT, P.G., ELLIOT, S.F., GRAY, L.J., HEWITSON, T.W., LECHMERE, D.J., LEWIS, D. & NEEDLEMAN, P.D. (1996). An alternative to the net premium valuation method for statutory reporting. *B.A.J.* **2**, *forthcoming*.

SHARPE, W. (1964). Capital asset prices. *Journal of Finance*, **19**, 425-442.

SHERRIS, M. (1992). Portfolio selection and matching: a synthesis. *J.I.A.* **119**, 87-105.

THORNTON, P.N. & WILSON, A.F. (1992). A realistic approach to pension funding *J.I.A.* **119**, 229-312.

WALTER, C. (1990). Mise en évidence de distributions Lévy-stables et d'une structure fractable sur le marché de Paris. *Transactions of the 1st AFIR International Colloquium*, Paris, **3**, 241-259.

WHALLEY, A.E. & WILMOTT, P. (1994) *.An asymptotic analysis of the Davis, Panas and Zariphopoulou model for option pricing with transaction costs.* OCIAM working paper, Mathematical Institute, Oxford.

WILKIE, A.D. (1986). Some applications of stochastic investment models. *J.I.A.S.S.* **29**, 25–50.

WILKIE, A.D. (1987). An option pricing approach to bonus policy. *J.I.A.* **114**, 21-77.

WILKIE, A.D. (1995). More on a stochastic asset model for actuarial use. *B.A.J.* **1**, 777-964.

WISE, A.J. (1984). The matching of assets to liabilities. *J.I.A.* **111**, 445-485.

WISE, A.J. (1987). Matching and portfolio selection, Part 2. *J.I.A.* **114**, 551-568.

APPENDIX A

TECHNIQUES FOR DYNAMIC OPTIMISATION

A.1 *The Optimisation Problem*

A.1.1 Dynamic optimisation techniques are almost absent from the actuarial literature. I hope that the examples in Section 3 have convinced the sceptic that the subject is worth a second look. I also want to dispel the common myth that dynamic optimisation only works in contrived examination examples. One early reference is Benjamin (1984), who identified some dynamic methods in actuarial work with established principles in control theory. More recently, Sherris (1992) wrote down some generic equations to be solved in dynamic asset-liability modelling, but to date, a good solution algorithm has been lacking. Most of the pioneering mathematics in this area has been developed by control engineers, and applied by financial economists, such as Merton (1971). The technique of dynamic optimisation can reasonably, therefore, be described as a technique of financial economics. In what follows we take much of the philosophy of Wise (1984). We are interested in a series of cash flows, at times t=0,1,2, ... T. We denote the cash flow at time t by C_t. These form a vector (C_0, C_1, C_2, .. C_T), which we denote in bold type as $\mathbf{C}$.

A.1.2 The cash flow vector of interest will vary according to the problem under consideration. In the first applications to defined benefit pension plans, the cash flow considered was the income from the assets plus contributions less benefit outgo. In insurance applications, we would be interested in transfers to shareholders' funds or in actual dividends paid. In developing these techniques, I have, therefore, started with the bare essentials, which means cash flow models. I have no preconceived notion of value, but we will see that, in the course of optimising cash flow models, some value concepts turn out to be useful intermediate steps.

A.1.3 The components of the cash flow vectors are not constants, but instead are random variables. Furthermore, there is an implied information structure that the cash flow C_t must be a function only of experience which has emerged by time t. In other words, it is not possible to look ahead, for example distributing at time 2 the surplus arising at time 3. Probability theorists would say that the vector $\mathbf{C}$ is *adapted*, while engineers would use the synonymous term *non-anticipating*.

A.1.4 We suppose that our client can specify a rule by which one cash flow vector is preferred to another. We then define a function U such that $\mathbf{C}$ is preferred to $\mathbf{C}'$ if and only if $U(\mathbf{C}) > U(\mathbf{C}')$. The quantity $U(\mathbf{C})$ is not a random number, but a known constant, so that random cash flows can be compared at the start. We emphasise that U is not a function of $\mathbf{C}$ which is calculated for each outcome separately, but rather a function of all outcomes at once, with their associated probabilities. For example, U might be an expected utility, but many other formulations are possible.

A.2 *The Bellman Approach*

A.2.1 The textbook approach to dynamic optimisation follows the original work of Bellman (1957). The idea is to break down a problem with many cash flows into a series of single time horizon problems. This technique only works analytically in rather special cases, but we can often use it to get approximate solutions.

A.2.2 Rather than reproduce theory which has been covered in depth elsewhere, I am content to demonstrate the Bellman approach by example. We consider a matching problem to pay a fixed liability L_T at some point T in the future. The investment objective is then to find an asset strategy to produce a final asset value A_T to minimise the expected squared surplus $\mathbf{E}[(A_T - L_T)^2]$.

A.2.3 We assume that the one-year mean-variance efficient frontier looks the same each year. The return factor (that is, $1+i$ in traditional notation) on the minimum risk portfolio has mean μ and standard deviation σ. If a portfolio has mean return different from μ, then we assume that a higher minimum variance applies, according to the formula:

$$\mathbf{Var}[\text{ return factor }] \geq \sigma^2 + \frac{1-\lambda}{\lambda} \{ \mathbf{E}[\text{ return factor }] - \mu \}^2.$$

Here, the parameter λ, with $0 < \lambda < 1$, is an indication of how much extra return might be obtained for a given degree of additional risk, larger values of λ corresponding to a larger reward for risk. This parabolic shape is roughly typical for many asset models, although few would fit it exactly.

A.2.4 Let us consider the situation at time $T-1$, when we have assets with value V_{T-1}. These are to be invested in such a way as to minimise the expectation $\mathbf{E}_{T-1}[(A_T - L_T)^2]$ subject to the efficient frontier constraints above. It is easily seen that the optimal investment strategy is to choose a portfolio with:

$$\mathbf{E}[\text{return factor}] = \lambda \frac{L_T}{V_{T-1}} + (1-\lambda)\mu$$

that is, a weighted average of the minimum risk return and the return required to meet the liabilities. The attained optimum is:

$$\mathbf{E}_{T-1}[(V_T - L_T)^2] = (1 - \lambda)(\mu V_{T-1} - L_T)^2 + \sigma^2 V_{T-1}^2.$$

We now define L_{T-1} by the relationship: $L_t = \left(\mu + \frac{\sigma^2}{(1-\lambda)\mu} \right) L_{T-1}$, and the attained optimum has the form:

$$\mathbf{E}_{T-1}[(V_T - L_T)^2] = \text{constant}_1 * (V_{T-1} - L_{T-1})^2 + \text{constant}_2.$$

A.2.5 At time T–2 we again have an optimisation problem, but since we know the optimum behaviour from T–1 to T, it now remains only to optimise the behaviour between T–2 and T–1. This involves minimising the expected value of (constant$_1$ * $(V_{T-1}-L_{T-1})^2$ + constant$_2$), which is equivalent to minimising the expected value of $(V_{T-1}-L_{T-1})^2$. However, we have already solved a problem of this form, from time T–1 to T, so we can apply the same solution as we had before. And so, using this process, we can work backwards inductively, requiring only one-year optimisation, at each time t–1 trying to hit a target L_t given by:

$$L_t = \left(\mu + \frac{\sigma^2}{(1-\lambda)\mu}\right)^{-(T-t)} L_T.$$

We can interpret L_t as a present value of L_T discounted at a return factor of $\mu + \frac{\sigma^2}{(1-\lambda)\mu}$ p.a. We can see how reserving on a particular basis arises naturally in the optimisation of asset-liability strategies. Indeed, we can now claim that this particular method of valuation and discount rate is more useful than any other method which might conceivably be devised, within the model and objectives we have set, because repeated pursuit of short-term goals based on this valuation basis leads to optimal behaviour over a longer term. In the case where σ=0 there is a risk-free asset, and the liabilities are discounted at market rates. For other values of σ there is no traded match for the liabilities, and so the market value concept does not immediately apply. The above formula is an example of *marginal value*; we see in Section 6 that marginal value is, in some sense, an extension of the market value

A.2.6 A further application of these optimisation techniques could be to justify some of the actuarial methods which may seem arcane to outside observers. In Section 3 we saw how a form of bonus reserve valuation may be justified in this manner. I have tried, without success, to produce similar justifications for other actuarial techniques, such as the assessed value of assets in pension fund work, net premium reserves in life assurance and equalisation reserves for non-life insurance. The justification need not necessarily be expressed in terms of shareholder value; the security of policyholder benefits, the desire of the regulators to avoid embarrassment and the plans of the board to further their own careers will all need to be taken into account. Each of these objectives and combinations, thereof, may be optimised using the methods of this appendix. A significant common thread is the importance of market values, even when the market is inefficient or distorted by frictional effects. My hunch is that traditional actuarial methods may be harder to justify, since the optimal methods for any of these objectives seem to be more deeply rooted in market prices than the actuarial methods are. I would like to be proved wrong in this regard, and encourage others to join the search. However, if I am right, then the profession ought to move to more market-based methods in line with the rest of the financial community. Dyson & Exley (1995) have outlined some of the practical problems which arise when using off-market asset values in pension fund work. Scott *et al.* (1996) have made similar observations regarding the net premium valuation method.

A.3 *Market Prices and Deflators*

A.3.1 Let us consider a cash flow vector $\boldsymbol{C}$. We might consider trading a certificate in the market carrying an entitlement to these cash flows. If the certificate is traded at time t, then only the cash flows at time t+1 and later are transferred. Let us suppose, at time t, that we can purchase the certificate for $A_t(\boldsymbol{C})$, the ask price. Similarly, we can sell the certificate for $B_t(\boldsymbol{C})$, the bid price. Both the cash flows and the transaction prices are defined net of any applicable taxes. We must have $A_t(\boldsymbol{C}) \geq B_t(\boldsymbol{C})$, the difference being transaction costs.

A.3.2 If the certificate cannot be purchased at time t, then we define $A_t(\boldsymbol{C}) = +\infty$. Similarly, if it cannot be sold we write $B_t(\boldsymbol{C}) = -\infty$. If the certificate trades costlessly, then we have $A_t(\boldsymbol{C}) = B_t(\boldsymbol{C})$. Thus, the framework can cope with complete absence of trading, frictionless trading or the intermediate situation which is trading with transaction costs. Since a negative sale is simply a purchase and vice versa, we have $A_t(-\boldsymbol{C}) = -B_t(\boldsymbol{C})$ and $B_t(-\boldsymbol{C}) = -A_t(\boldsymbol{C})$. Finally, we assume transaction costs are proportional to the transaction size, so that for $\lambda \geq 0$ we have $A_t(\lambda \boldsymbol{C}) = \lambda A_t(\boldsymbol{C})$ and $B_t(\lambda \boldsymbol{C}) = \lambda B_t(\boldsymbol{C})$.

A.3.4 We assume that the price does not depend on the current quantity held. In particular, it is quite possible to sell assets you do not own, so there are no restrictions on borrowing or short selling. The limitation can be relaxed somewhat if we consider only small deviations about an initial position. In this situation, assets which form part of the initial portfolio would be assigned the market bid and ask prices. Assets which are not held would be assigned a market ask price, but a bid price of $-\infty$, thus making short sales infinitely costly.

A.3.5 Let $\boldsymbol{D} = (D_0, D_1, D_2, .. D_T)$ be an adapted vector. We say that $\boldsymbol{D}$ is a *deflator* if for every cash flow vector $\boldsymbol{C}$ and each t we have:

$$D_t A_t(\boldsymbol{C}) \geq \mathbf{E}_t[\, D_{t+1}C_{t+1} + D_{t+2}C_{t+2} + \ldots + D_T C_T \,] \geq D_t B_t(\boldsymbol{C}).$$

Intuitively, we can think of a deflator as a stochastic generalisation of a discount factor. For example, let us imagine what would happen if $D_t = (1+i)^{-t}$ for some constant rate of interest i. In that case, the above condition says that the net present value of expected cash flows always lies within the bid-ask spread. In reality, such a simple relationship is not likely to hold with the same i for all cash flow vectors simultaneously, because we would expect to see some adjustment for risk in the discount rate.

A.3.6 It is not obvious that any deflators should exist at all. We shall see that, in fact, for most asset models, we can find at least one deflator. In particular, the existence of a deflator does not rely on any notion of market efficiency. If a model has wide trading spreads, then there will be many possible deflators, since there is a lot of scope to squeeze different discounting formulae between the bid and ask prices. If nothing trades at all, then any positive vector is a deflator. At the other extreme, if every cash flow trades frictionlessly, then the market is said to be *complete,* and there is a unique deflator up to multiplication by a constant scalar.

A.3.7 We will see that identification of the set of deflators plays an important part in dynamic optimisation. For some stochastic models, such as random walk models or my jump equilibrium model, it is relatively easy to identify the set of possible deflators. In other cases, such as the Wilkie model, the algebra is much more difficult. It follows that dynamic optimisation will be easier in the former case and more difficult in the latter.

A.4 *The Differential Approach*

A.4.1 Traditionally, one would hope to solve optimisation problems simply by differentiating and setting to zero. The Bellman approach looks much more complex than is really necessary. I have found that the differential approach is often easier, and is also more general than the Bellman approach. The differential approach is particularly helpful when the market is complete, or nearly complete, in the sense that the set of deflators is small. The reason for the popularity of the Bellman approach in the literature seems to be mainly historical.

A.4.2 We consider again the problem, in ¶4.1.4, of maximising $U(\boldsymbol{C})$. We suppose that assets providing a cash flow vector $\boldsymbol{F}$ are already in place, and that a revenue r has just been received, of which some may be invested to give cash flows $\boldsymbol{I}$. The problem is then to select $\boldsymbol{I}$ solving:

$$\max U(\boldsymbol{F} + \boldsymbol{I}) \text{ subject to } I_0 + A_0(\boldsymbol{I}) \leq r.$$

The inequality on the right is sometimes called the *budget constraint.* By redefining U, we may assume that $\boldsymbol{F} = 0$, so that the problem now becomes:

$$\max U(\boldsymbol{C}) \text{ subject to } C_0 + A_0(\boldsymbol{C}) \leq r.$$

A.4.3 We now define a gradient concept, extending that of classical vector calculus. Let us consider a fixed $\boldsymbol{C}$, and let us suppose we can find an adapted vector $\boldsymbol{G}$ such that for any adapted vector $\boldsymbol{H}$ we have, for small ε:

$$U(\boldsymbol{C}+ \varepsilon \boldsymbol{H}) = U(\boldsymbol{C}) + \varepsilon \mathbf{E}[G_0H_0 + G_1H_1 + G_2H_2 + .. G_TH_T] + o(\varepsilon).$$

Then we say that $\boldsymbol{G}$ is the *gradient* of U at $\boldsymbol{C}$, writing $\boldsymbol{G} = U'(\boldsymbol{C})$. Under regularity assumptions on U which are usually satisfied in practice, we can find such a $\boldsymbol{G}$, in which case we say that U is differentiable.

A.4.4 We can now consider the *marginal value* M_t at time t of a series of cash flows H. This is the cash flow at time t, which one could substitute for all the original cash flows subsequent to t in any event F observable at time t. To first order, this condition is:

$$\mathbf{E}[G_t M_t ; F] = \mathbf{E}[G_{t+1} H_{t+1} + .. + G_T H_T ; F]$$

which implies that the marginal value is given by the conditional expectation:

$$G_t M_t = \mathbf{E}_t[G_{t+1} H_{t+1} + .. + G_T H_T].$$

We note that this is a positive linear form, as required of value concepts in ¶2.5.2.

A.4.5 We can now derive a first order necessary condition for optimality of a cash flow vector $\boldsymbol{C}$. Let us suppose that $\boldsymbol{C}$ is optimal, and let us consider another vector $\boldsymbol{H}$ of cash flows. We consider the opportunity to purchase or to sell at time t, a quantity ε of the cash flows in $\boldsymbol{H}$, starting from time t+1. By hypothesis on $\boldsymbol{C}$, the resulting cash flows will not be preferred to the original situation. Expanding to first order in ε, this gives the inequalities:

$$G_t A_t(\boldsymbol{H}) \geq \mathbf{E}_t[\ G_{t+1} H_{t+1} + G_{t+2} H_{t+2} + \ldots + G_T H_T\ ] \geq G_t B_t(\boldsymbol{H})$$

where $\boldsymbol{G}$ is the gradient of U at $\boldsymbol{C}$. These inequalities are exactly the condition for $\boldsymbol{G}$ to be a deflator in the sense of ¶4.2.5. Thus, the first order necessary condition for optimality of $\boldsymbol{C}$ is that $U'(\boldsymbol{C})$ is a deflator. Alternatively, this can be interpreted as saying that the marginal value of any asset lies within the market bid-ask spread. This has a natural interpretation as follows. Let us suppose, for a contradiction, that the marginal value exceeded the market ask price. Then the investor could improve his utility by lifting the market ask, and so the original asset allocation was not optimal. As the investor continues to buy, he accumulates a concentration of risk which reduces the marginal value of the purchased cash flows. This will continue until the marginal value falls within the bid-ask spread, at which point an optimum is attained. Similar considerations apply on the bid side.

A.4.6 When stating the original problem, we deliberately avoided reference to any form of valuation. However, in the mathematical solution, we see that three value concepts — bid, ask and marginal — turn out to be important. Market values arise naturally as part of the mathematical solution to an optimisation problem, rather than any preconceived dogma as to what value concepts are important. This result holds even if markets are inefficient or subject to frictional costs; factors which are often cited, in ignorance, as good cause for abandoning market value in favour of an 'assessed value' approach.

A.4.7 One striking feature of this result is the common form of the optimality criterion for many different types of investor. It is this that makes it possible to construct economic models where each agent is solving his or her own optimisation problem, even though the underlying objective functions are diverse, with different liabilities to match, different risk tolerances and different time horizons. The commonality arises, not because of similar preferences, but because of similar budget constraints. The budget constraints are similar because everyone sees the same prices. This elegance (and hence the ease of model construction) disappears when different participants are subject to different constraints, for example because of asymmetries in fiscal treatment. In my modelling experience, I have found it easiest to ignore these complicating factors when building the underlying economic model, reintroducing them at the end in the context of my particular client.

A.4.8 In the case of a complete market, in the sense of ¶A.3.6, this first order condition determines $U'(C)$ uniquely up to a constant multiple. If the multiple is known, then we can work backwards to find the cash flow $\boldsymbol{C}$. The appropriate multiple is determined by the budget constraint. On the other hand, if the market is not complete, then it is harder to work out $U'(C)$, because the deflator is not unique. There is a method for finding it which works most of the time, and is outlined below in ¶A.5.4.

A.5 *The Dual Optimisation Problem*

A.5.1 In some situations an analytical dynamic optimisation may not be practical, because of time or budget constraints, or because the mathematics is intractable. An experienced modeller may be able to guess at a reasonably good dynamic strategy, and determine how good the strategy is by means of simulations. As a benchmark, it is useful to have an upper bound on how good a strategy could possibly be.

A.5.2 In order to derive such upper bounds, it is helpful to define the *Fenchel conjugate* U^* of the function U. For an adapted vector $\boldsymbol{V}$, we define $U^*(V)$ as:

$$U^*(V) = \inf_C \mathbf{E}[C_0V_0 + C_1V_1 + C_2V_2 + \dots C_TV_T] - U(C).$$

The infimum is taken over all cash flow vectors $\boldsymbol{C}$, not just those which are feasible for a particular application.

A.5.3 We can now develop an upper bound on $U(C)$, subject to the constraint $C_0 + A_0(C) \leq r$. For any deflator $\boldsymbol{D}$, it is easy to prove the upper bound:

$$U(C) \leq D_0 r - U^*(\boldsymbol{D}).$$

A.5.4 It is instructive to consider the deflator $\boldsymbol{D}$ which produces the best (that is, the lowest) upper bound for $U(C)$. We can therefore consider the problem:

$$\min_D \; D_0 r - U^*(\boldsymbol{D}) \text{ subject to } \boldsymbol{D} \text{ being a deflator.}$$

This is sometimes called the *dual problem*, relative to the original *primal problem* which was:

$$\max_C U(\boldsymbol{C}) \text{ subject to } C_0 + A_0(\boldsymbol{C}) \le r.$$

Under certain continuity and convexity conditions, the primal and dual problems attain the same optimum. The importance of this result in financial modelling was established by Benjamin (1959), who applied the theory outlined here to interest rates. We note further that the optimal value of $\boldsymbol{D}$ for the dual problem is exactly the gradient $U'(\boldsymbol{C})$ at the optimum of the primal problem. If there are no deflators, then the dual problem is infeasible, so we can obtain no upper bound on the primal problem. This makes perfect sense; in the absence of a deflator one would expect to find arbitrage opportunities, which implies infinite utility.

A.5.5 As an example, let us consider the power law formulation, where, for some $\kappa > 1$ and positive weight vector $\boldsymbol{W}$, we have:

$$U(\boldsymbol{C}) = -\frac{1}{\kappa - 1}\mathbf{E}\sum_{t=0}^{T} W_t^{\kappa} C_t^{-\kappa+1}.$$

After a few lines of algebra, we can readily find $U^*(\boldsymbol{V}) = \frac{\kappa}{\kappa - 1}\mathbf{E}\sum_{t=0}^{T} W_t V_t^{1-\frac{1}{\kappa}}$. We use this result in Section 5 to calibrate a stochastic economic model.

A.5.6 We can use this dual problem to obtain some insight into the shape of mean-variance efficient frontiers. Let us denote by (σ, μ) the standard deviation and mean of the final cash flow C_T for an invested portfolio, resulting from its realisation at time T. We denote the annual premiums by $p_0, p_1, \ldots, p_{T-1}$; in general these may be stochastic. We then consider the problem of maximising $U(\boldsymbol{C})$, defined for some constant a by:

$$U(\boldsymbol{C}) = \begin{cases} -\mathbf{E}\{(a - C_T)^2\} & C_0 = -p_0;\ C_1 = -p_1; \ldots C_{T-1} = -p_{T-1} \\ -\infty & \text{otherwise.} \end{cases}$$

We can calculate the Fenchel conjugate of this U, which is:

$$U^*(\boldsymbol{V}) = a^2 - \mathbf{E}\left\{\left(\frac{V_T}{2} - a\right)^2\right\} - \mathbf{E}\{p_0 V_0 + \ldots + p_{T-1} V_{T-1}\}.$$

Now, using the upper bound for $U(\boldsymbol{C})$, and using the fact that the original cash holding $r=0$, we can see that for any deflator $\boldsymbol{D}$:

$$\sigma^2 + \mu^2 - 2a\mu + a^2 \ge a^2 - \mathbf{E}\left\{\left(\frac{D_T}{2} - a\right)^2\right\} - \mathbf{E}\{p_0 D_0 + \ldots + p_{T-1} D_{T-1}\}$$

and furthermore, since for any $\lambda > 0$, $\lambda\boldsymbol{D}$ is also a deflator, we have:

$$\sigma^2 + \mu^2 - 2a\mu + a^2 \geq a^2 - \mathbf{E}\left\{\left(\frac{\lambda D_T}{2} - a\right)^2\right\} - \lambda \mathbf{E}\{p_0 D_0 + \ldots + p_{T-1} D_{T-1}\}.$$

For a given value of μ, we can, by differentiation, determine the values of a and λ which give the tightest bounds for σ. The result is:

$$\sigma \geq \frac{\left|\mu \mathbf{E}(D_T) - \mathbf{E}(p_0 D_0 + \ldots + p_{T-1} D_{T-1})\right|}{\text{standard deviation } (D_T)}.$$

These are the inequalities I used to determine upper limits on efficient frontiers in Section 3.

A.6 *Existence of Risk-Neutral Probabilities*

A.6.1 We can now tie up some loose ends regarding the link between deflators and risk-neutral laws. We take a simple model with a single time period from 0 to 1. The uncertainty during the period is generated by a random variable X with density function $f(x)$. We take a deflator $\boldsymbol{D} = (D_0, D_1)$ where D_0 is a constant, and D_1 is a function of X. We define an interest rate i, called the *risk-free rate*, by $i = \frac{D_0}{\mathbf{E}(D_1)} - 1$. We then define a *risk-neutral* probability density function $f_{RN}(x)$ by $f_{RN}(x) = (1+i)D_1(x)f(x)/D_0$. This is easily seen to be positive and integrate to 1, so is a probability density function. We notice, in passing, that there is no need for a risk-free asset to exist in the market for this algebra to work. If there is a risk-free asset then it will earn the risk-free rate, hence the terminology. In other circumstances the risk-free rate is shorthand for the expression above, and does not equate directly to any rate quoted in the market. There is thus a one-to-one correspondence between deflators and risk-neutral laws.

A.6.2 We have, by definition of deflator, for any cash flow $\boldsymbol{C}$:

$$A_0(\boldsymbol{C}) \geq \frac{1}{1+i} \mathbf{E}_{RN}(C_1) \geq B_0(\boldsymbol{C})$$

where $\mathbf{E}_{RN}$ is the expected value under the risk-neutral probability density function. We can see that the discounted expected value, under the risk-neutral probability, lies between the bid and ask prices. A similar result extends to multi-period settings. There is, thus, a close connection between the existence of risk-neutral laws and the existence of deflators. We can see that the risk-neutral law gives more weight to scenarios where D_1 is high and less to where D_1 is low. We can identify the former with pain and the latter with pleasure. The interpretation is that it is cheap to buy assets which pay out in pleasant scenarios, but more expensive to buy insurance against unpleasant scenarios. This explains why a risk premium is available for taking market risk — the payoff is then greater in the most pleasant scenarios.

A.6.3 If the cash flow C is traded frictionlessly, then we can calculate the return R_p obtained on investing to obtain C_1. It is easy to see that it satisfies the relationship:

$$\mathbf{E}(R_p) = i - (1+i)\mathbf{Cov}\left(R_p, \frac{D_1}{D_0}\right)$$

so that the expected return is the risk-free rate minus a covariance adjustment. An investor can, therefore, earn extra expected returns above the risk-free rate by choosing a portfolio which is negatively correlated with D_1/D_0. We can now identify precisely those trading positions which may generate extra expected returns and those which would not. This is the essence of arbitrage pricing theory as developed by Ross (1976). I emphasise that, to get to this point, we have not needed to assume any form of market efficiency.

A.6.4 Readers may notice a remarkable similarity of ¶A.6.3 to the CAPM. However, the CAPM goes further than ¶A.6.3, and claims that the market return is a linear decreasing function of D_1/D_0. This means that positive correlation with market risk is rewarded with extra expected returns, but other risks are not. Most derivations of the CAPM make questionable assumptions regarding either normality of investment returns or quadratic utility functions. For other plausible return distributions and utility functions, the market return is a convex decreasing function of D_1/D_0. Very often, this subtle correction to the CAPM will make little or no numerical difference. In other words, although some of the assumptions underlying the CAPM are questionable, the model is robust to deviations from those assumptions. I have suggested reasons for this robustness in Section 6.

A.7 *Portfolio Optimisation with Option Prices*

A.7.1 One particularly interesting, and tractable, application of these optimisation techniques is to the inclusion of options within an asset portfolio. This has been considered by Lee (1993), and we now generalise his results within a Black-Scholes framework.

A.7.2 We consider a market with two assets — cash which earns a risk-free rate r, expressed as continuously compounded, and a share which pays dividends continuously at a yield q. The capital value of the share is denoted by S. It is assumed that log S performs a geometric random walk with volatility σ, so that, for some parameter $\alpha \geq 0$, we have, for $t<u$:

$$\log(S_u/S_t) \sim \mathrm{N}[(r-q+(\alpha-\tfrac{1}{2})\sigma^2)(u-t)\ ,\ \sigma^2(u-t)\ ].$$

A.7.3 The larger the value of α, the greater the reward for risk. If $\alpha=0$, then the expected total return on equity is exactly that of cash.

A.7.4 Options are priced in accordance with the dividend adjusted Black-Scholes formula. Considering European style options expiring at time u with strike K, the prices at time t are given as in Section 2.6. We notice that these formulae do not depend on α, the reward for risk, since the same risk-neutral probability law applies for all α.

A.7.5 We can identify a deflator D whose components are given by:

$$D_t = \exp[\ -\{(1-\alpha)r + \alpha q + \alpha(1-\alpha)\sigma^2/2\}t\ ]\ S_t^{-\alpha}\ .$$

This is the unique deflator (up to multiplication by a constant) for this model. As anticipated in ¶A.6.2, high values of the deflator occur in painful scenarios when equities have performed badly, while conversely, low values of the deflator occur in more pleasant scenarios.

A.7.6 This immediately enables us to determine the portfolios which are efficient from a mean-variance perspective. Let us consider investing for a time horizon T with utility function $u(C_T)$. Since the deflator is unique up to a constant multiple, the first order condition is:

$$u'(C_T) = \lambda D_T$$

so that $C_T = [u']^{-1}(\lambda D_T)$. The constant λ is chosen so that the initial budget constraint is satisfied. Using such techniques we can, for example, deduce that all mean-variance efficient portfolios have a final cash flow of the form $C_T = a - bS_T^{-\alpha}$ for constants a and b. This solves the problem proposed by Lee (1993). Interestingly, this would imply that all investors are net sellers of options. The dynamic hedge for such an option position would involve decreasing exposure to risky assets as markets rise. It was precisely such strategies which we found to be on the efficient frontier in Section 3. Since the net supply of options must be zero, it follows that some investors, namely the option buyers, must hold portfolios which are not mean-variance efficient.

A.8 *Alternative Formulations of Investor Preferences*

A.8.1 The actuarial literature contains a fair amount of criticism of utility theory, for example Clarkson (1995), who hints that financial economists may have been rather inflexible in this matter. Various other criteria have been suggested.

A.8.2 It has been proposed that investors might want to maximise expected return, subject to a limit on the probability of getting one's money back. The maximum is achieved by holding a zero coupon bond to provide a money-back guarantee. The remainder of the money is spent on call options. The higher the strike of the call options, the higher the expected return from the whole arrangement. There is no upper bound to the expected return. Although the objective function seems reasonable, the answer obtained does not, in my view, offer any useful insight to the fund manager. It is easily demonstrated that this result is not specific to the Black-Scholes world, but applies more generally to any model with a discount which is not bounded above. Virtually all models, including my own stochastic model, satisfy this condition. I note, in passing, that the same problems arise when solving the mathematically very similar problem of maximising some percentile of the return distribution.

A.8.3 Investors may wish to trade off the mean μ of the cash flow C against the semi-variance, where the semi-variance is defined as $\mathbf{E}[\max\{\mu-C, 0\}^2]$. This behaves somewhat like the variance, but only captures downward deviations from the mean. In the notation of ¶A.1.4 and using Lagrange's principle, the efficient portfolios will maximise objectives $U(\mathbf{C})$ of the form:

$$U(C_0,\ C_1)=\begin{cases}-\infty & C_0<0\\ \mathbf{E}(C_1)-\lambda\mathbf{E}\left[\max\{\mathbf{E}(C_1)-C_1,0\}^2\right] & C_0\geq 0\end{cases}$$

where different $\lambda > 0$ parameterise the efficient frontier. In order to perform the optimisation we calculate the Fenchel conjugate $U^*(\mathbf{V})$. For most vectors $\mathbf{V}$ we have $U^*(\mathbf{V}) = +\infty$, the exception being when $V_0 \geq 0$, $\mathbf{E}(V_1) = 1$ and V_1 is bounded below by some constant m. In this case, the Fenchel conjugate is readily shown to be $-\frac{\mathbf{E}[(V_1-m)^2]}{4\lambda}$. Three cases then occur. If (case I) there is no deflator $\mathbf{D}$ for which $U^*(\mathbf{D})$ is finite, then the problem is unbounded, and infinite expected returns can be obtained with finite downside risk. If for some deflator $\mathbf{D}$ we have finite $U^*(\mathbf{D})$, then there are two more cases, according to whether or not there is a strictly positive probability that D_1 attains its minimum value. If so (case II), then the optimum is attained, but not uniquely, so that the portfolio selection problem is still unsolved. Otherwise (case III, which includes the Black-Scholes model), the optimum is not attained, but one can get arbitrarily close using bond plus call options, as in ¶A.8.2. In each of these cases the optimisation has shed little light on the asset allocation decision. Instead, we have discovered that the portfolio selection criterion was not as robust as some authors would have us believe.

A.8.4 The objectives in the preceding sections were developed to remedy some perceived deficiencies of the utility approach for choosing portfolios when a small number of investments were involved. Empirically, the methods may seem to give plausible answers when applied in such situations. It is the introduction of derivatives, and indeed infinitely many possible derivatives, which seems to cause problems. In such situations the utility approach, despite its limitations, does seem to be more robust than alternative methodologies. The reason for the continued popularity of utility functions among economists is not intransigence, but the failure of any of the alternatives to be useful for solving practical problems.

APPENDIX B

METHODS FOR CALIBRATING MODELS

B.1 *Models to be Calibrated*

B.1.1 This section considers the calibration of the economic models described in this paper. This is not strictly an application of financial economics, but rather of econometrics.

B.1.2 In this section I have calibrated the random walk model, the Dyson & Exley model and my own model. I have not calibrated the chaotic and fractal models separately; instead I have taken the parameters for the random walk model and substituted different series for the error terms. I have not attempted to recalibrate the Wilkie model. I have followed Wilkie (1995), with the exception of the property series, where I have used Daykin & Hey (1990).

B.1.3 In each case I have calibrated the model in two steps. The first step is to examine the historic variance and covariance (and higher moments) structure of the economy. The second step is to apply financial economics to determine the appropriate return for each asset class, given the risks. I have used this to estimate means, instead of the more commonly employed historic experienced return.

B.1.4 I have made no consideration of government policy or fundamental analysis. For the Dyson & Exley model, and also for my own model, the initial term structures are inputs, so that any known policy changes are taken into account indirectly, to the extent that they are anticipated by the market.

B.2 *Statistical and Economic Calibration*

B.2.1 I have considered the following thought experiment. Let us suppose that an investor has a logarithmic utility function, and wishes to hold a portfolio containing U.K. equities and U.K. gilts. The optimal asset proportions will vary according to the relative expected returns on the two asset classes. Let us suppose, further, that in any year the annual outperformance of one relative to the other, in log terms, has a normal distribution with constant mean μ which is to be determined, and known standard deviation of $\sigma = 15\%$. We now consider various ways of estimating μ.

B.2.2 We consider the optimal asset mix for different values of μ. It can be shown that the optimal portfolio will contain a mixture of the two asset classes, provided that $|\mu| < 1.13\%$ ($= \sigma^2/2$), which we call the *interesting range*. Outside this range the optimal portfolio is invested 100% in whatever asset class has the higher expected return.

B.2.3 A purely statistical approach to the estimation of μ would be to use some historic data going back, say, 50 years. A two-sided 95% confidence interval for μ would then have a width of 4.1%, that is nearly twice as wide as the interesting range. More frequent sampling intervals would do nothing to reduce this standard error.

B.2.4 Incidentally, the 95% confidence interval for the estimated standard deviation is approximately from 12% to 18%. While this uncertainty still seems substantial, the uncertainty in portfolio selection is still dominated by uncertainty in mean returns. If monthly data were available, the 95% confidence interval for the standard deviation shrinks to (14.15%, 15.85%). However, this figure is highly sensitive to the normality assumption. For distributions other than normal, the confidence intervals tend to be much wider.

B.2.5 It seems to me that historic data alone cannot, even in principle, provide estimates of mean returns which are tight enough to be meaningful for asset allocation decisions. We are forced to look more deeply into the price formation mechanism in order to build more structure into the problem. Furthermore, tests of any assumed theory against historic return data will necessarily be very weak (that is, unlikely to refute the theory when it is wrong), since, if the data are inadequate for calibrating the initial model, they are also likely to be inadequate for testing the theory.

B.2.6 Various approaches to the estimation of expected returns have been suggested in Appendix B of Mehta (1992). The approaches suggested for equities are:

— an assessment of corporate growth prospects;
— a consideration of return on capital employed;
— an analysis of historic returns; and
— consensus forecasts.

I have used a different approach to estimation of the means. I model investment decisions using utility functions, and choose means such that the theoretical demand for different investment classes equates to the known supply.

B.3 *Data*

B.3.1 The data from which I have started are an inflation index and annual total return indices for cash, bonds, equities, index-linked gilts and property. Investment return data have been taken from the BZW equity gilt study (1995), and more recently from CAPS. I have used JLW property returns, as adapted by Blundell (1995).

B.3.2 My inflation, cash, equities and gilts series start with the return during 1919, and continue until 1994. My property series starts in 1970, because that is the earliest I could get. Index-linked gilts were first issued in 1981, so there are no data available before then.

B.3.3 For each of the models I have fitted, there is a transformation which, if the model is true, turns the investment series into independent identically distributed random vectors. For the random walk model, the relevant vector is:

$$Y_t = \begin{pmatrix} \log(1+\text{inflation}_t) - QA * \log(1+\text{inflation}_{t-1}) \\ \log(\text{equity real return factor}) \\ \log(\text{gilt real return factor}) \\ \log(\text{cash real return factor}) \\ \log(\text{property real return factor}) \\ \log(\text{index-linked gilt real return factor}) \end{pmatrix}.$$

The real return factor would be written in actuarial terms as $\frac{1+i}{1+e}$, where i is the nominal return achieved and e is the rate of inflation. For the cointegrated models, the relevant vector is:

$$Y_t = \begin{pmatrix} \log(1+\text{gilt return}) - \log(1+\text{cash return}) \\ \log(1+\text{equity return}) - \log(1+\text{cash return}) \\ \log(1+\text{property return}) - \log(1+\text{cash return}) \\ \log(1+\text{index-linked gilt return}) - \log(1+\text{cash return}) \end{pmatrix}.$$

It makes no difference whether these returns are measured in real or nominal terms, provided they are consistent.

B.4 *Historic Variances and Covariances*

B.4.1 For each of these models we have a series of terms Y_t, which are supposedly independent from one time to the next. In each model these are to be expressed as a linear function of a vector whose components are independent. Thus, we look for a square matrix L such that $Y_t = LX_t$, where X has independent components.

B.4.2 We first solve a rather simpler problem of finding a matrix L such that X, as defined above, has uncorrelated components. If the components are independent, then they are also uncorrelated, but not conversely. Without loss of generality, we can also assume that the components of X have unit variance. The variance-covariance matrix of Y is then LL^T, where the superscript T denotes a matrix transpose. The calibration of the models involves two steps — firstly the estimation of a historic variance-covariance matrix and secondly the extraction of a form of matrix square root.

B.4.3 The estimation of historic variances and covariances ought to be simple, using standard formulae. The problem is that, for earlier time periods, some of the vectors are incomplete because of missing information on property and index-linked gilts. The naive solution is to calculate the variances and covariances over the longest possible time, that is, using all information for which the relevant components are defined. This means that the variance of gilt returns would be measured over a longer period than the variance of index-linked gilts or the covariance of the two. If the true model were as proposed, such a calibration would be an optimal procedure.

B.4.4 Difficulties arise because we are not fitting the correct model. Instead, we are fitting models which are easy to implement. The real world is much more complex than the models we are trying to fit. One particular pattern which none of the models captures is the ARCH effects, which describe the tendency of the economy to alternate between calm and volatile periods. For example, the period since 1980 has seen stable levels of inflation by comparison with the rest of the century. It may not be appropriate to calibrate an inflation model based only on this stable period without some reference also to much greater inflation volatility in previous decades. However, when calibrating index-linked gilts, we have no data prior to 1980, so we are forced to use data only from a period of inflationary stability. This produces an inconsistent model, where inflation volatility is calibrated from a long and turbulent history, while index-linked gilts are based on a stable period. It is reasonable to suppose that, if inflation were, in future, to be as volatile as the past century, then index-linked gilts are likely to be far more volatile than they have been in the last fifteen years. The obvious estimation procedure, therefore, understates the volatility of index-linked gilts. This observation has important consequences. Many asset-liability studies, particularly for pension funds, recommend a significant strategic allocation in index-linked gilts, such that, if the results for all funds were aggregated, the quantity of index-linked gilts required would exceed those in issue by a substantial multiple. This anomaly is largely due to the understatement of index-linked gilt volatility in carelessly built stochastic models.

B.4.5 One solution to the ARCH problem is to fit an ARCH model. However, other aspects of the models would then become intractable. I have carried out some empirical tests using ARCH models and their simpler non-ARCH versions to see whether the ARCH effect has a significant bearing on optimal actuarial decisions, and, in the main, I found that it did not, so I have not followed the ARCH route. Instead, I have made some simple adjustments to the data series to adjust for the ARCH effect, and then fitted simpler models. My adjustment involves splitting the data into sub-periods, such that, within each sub-period, the same data are available each year. This gives three sub-periods: 1919-69 for which I have inflation, short-term interest rates, gilt and equity total returns; 1970-81 for which I also have property total returns and 1982-94 for which I also have index-linked gilt total returns. I calculate standard deviations for each data series and sub-period. I then express these volatilities approximately as a product of one factor relating to the asset class and another factor relating to the time period, by fitting a linear model to log standard deviations. The period-dependent factor gives me an indication of the relative volatilities of the periods covered. I use these relative volatilities as scaling factors to the original residuals, to obtain a series of adjusted residuals with the property that, for any asset, the volatility is more or less the same for each sub-period. This adjustment overcomes the ARCH effect described above, so I can now apply traditional methods for calculating a variance-covariance matrix V. I then determine the Cholesky square root, which is the unique lower triangular matrix L with non-negative diagonal elements such that $V = LL^T$. This whole process is rehearsed twice; once in six dimensions for the random walk and once in four dimensions, which covers the cointegrated models. Full working is provided in the spreadsheets which I have deposited on the Internet.

B.4.6 For the random walk model I obtained the following matrix L:

Inflation residual	0.0560	0.0000	0.0000	0.0000	0.0000	0.0000
Equities	−0.0529	0.2196	0.0000	0.0000	0.0000	0.0000
Gilts	−0.0716	0.0702	0.1045	0.0000	0.0000	0.0000
Cash	−0.0406	0.0016	0.0058	0.0332	0.0000	0.0000
Property	−0.0488	0.0743	−0.0676	−0.0619	0.2056	0.0000
Index-linked gilts	−0.0164	0.0617	0.0952	−0.0006	0.0049	0.0558

while for the cointegrated models I have the L-matrix for returns relative to cash:

Gilts	0.1113	0.0000	0.0000	0.0000
Equities	0.0890	0.1704	0.0000	0.0000
Property	0.0084	0.0426	0.1402	0.0000
Index-linked	0.0371	0.0037	0.0051	0.0611

B.4.7 The astute reader will have smelled something fishy at this point. I could have taken the variance-covariance matrix for the random walk model and deduced the variance covariance matrix for the other two models using linear algebra. If I do so, the answers are not consistent, because of the adjustments I have made to allow for ARCH effects. The L matrix I would have obtained via the random walk route is:

Gilts	0.1285	0.0000	0.0000	0.0000
Equities	0.1235	0.1833	0.0000	0.0000
Property	0.0090	0.1005	0.2277	0.0000
Index-linked	0.1037	0.0034	0.0025	0.0755

We can see that most of these entries are larger (and significantly so) than those in the matrix I actually used. This implies that, if I look at simulated returns relative to cash, the random walk model will show greater volatility than the other approaches, even though they are all calibrated from the same data set. The reason that this happens is because of the distinction between a 'long-run mean', which corresponds to the average of a long data series, and a 'contemporary mean', which is a short-term conditional expectation given recent history. In the cointegrated models, contemporary mean returns for a particular asset class are allowed to fluctuate over time, while the relativities between asset classes are more stable. By contrast, in the random walk model, the contemporary mean is assumed to be equal to the long-term mean. Any variability in the data due to long-term changes in contemporary means cannot be accommodated accurately within a random walk framework, and so ends up, by default, loaded into the short-term volatilities. This gives some annual volatilities for the random walk model which may seem implausibly high, such as 5.6% for the rate of inflation, or 5.3% for cash.

B.5 *Resolution into Independent Components*

B.5.1 We have expressed residual vectors Y in the form LX, where X is supposed to have independent components. Our method of construction ensures only that $L^{-1}Y$ has uncorrelated (but not necessarily independent) coefficients. Furthermore, there was some arbitrariness in the choice of L. If we had chosen a different L still satisfying $LL^T = V$ (of which there are infinitely many), then the new $L^{-1}Y$ would also have uncorrelated coefficients. We have not yet addressed the question of whether some choices of L result in vectors X whose coefficients are independent, while others merely produce components with zero correlation.

B.5.2 When models are based on multivariate normal distributions, these questions are easy to answer, because, in this framework, uncorrelated variables are independent. Therefore, the arbitrariness in the choice of L has no consequence; any choice will result in the same simulated model. This applies to the random walk model and also to the Dyson & Exley model. It is only when we deviate from this normal framework, as I do in my model, that the choice of L starts to matter. We now consider this model in more detail.

B.5.3 Unfortunately the standard method of maximum likelihood is unsuitable for fitting the shifted gamma distributions which arise in my model, since, by taking $\alpha<1$, the gamma density explodes at the origin, and, by shifting the origin, we can make the likelihood function explode at one point, irrespective of the values at the other sample points, and thus the likelihood maximisation problem is unbounded. Another way to tackle the problem is to choose a statistical test of independence and select the linear transformation for which the implied sample of X minimises the test statistic. Unfortunately most statistical tests of independence are actually tests of correlation, and therefore get us no further. The non-parametric tests which do truly measure independence take a discrete set of values, and so are of little use for optimising continuous variables. Failing the published methods, I had to develop a new methodology. I have applied this only to the first two components of Y (that is the returns on equity and gilts relative to cash), since the methodology is rather data intensive and is unlikely to be helpful for the property and index-linked time series, where only a short span of data is available.

B.5.4 The essence of my method is as follows. If the components of X are independent, then the moment generating function of X factorises as a product of the moment generating functions of each component. This implies a series of relationships of the form:

$$\mathbf{E}\{\exp(a_1X_1+a_2X_2)\}\mathbf{E}\{\exp(b_1X_1+b_2X_2)\} = \mathbf{E}\{\exp(a_1X_1+b_2X_2)\}\mathbf{E}\{\exp(b_1X_1+a_2X_2)\}.$$

Given samples of a vector X, we can test for independence by calculating each side empirically for a range of parameters (a_1, a_2, b_1, b_2). The extent to which the relationship is satisfied can then be observed. In my algorithm, I take various (a_1, a_2, b_1, b_2) and then measure the deviation from independence by the difference of the logs for each side of the above equation. Summing the squares of all these test variables gives a measure of the degree to which the components of X were independent.

B.5.5 I then look for the transformation L such that $X = L^{-1}Y$ minimises the sum of squares as in ¶B.5.4. The optimisation is fiendishly difficult — observe that the optimum ought to be invariant under changes in sign and permutations of the variables, so that the objective function will have 8 local minima. Furthermore, these local minima all lie in a curved valley which forms a hypersphere surrounding the origin, corresponding to vectors X with uncorrelated components, which would cause problems for most general purpose optimisation routines. On the bright side, the point at which the optimum is attained does not seem sensitive to the choices of (a_1, a_2, b_1, b_2). The matrix L which turns out to be optimal according to this criterion is:

Gilts	0.1083	0.0256	0.0000	0.0000
Equities	0.0474	0.1863	0.0000	0.0000
Property	−0.0016	0.0434	0.1402	0.0000
Index-linked	0.0353	0.0122	0.0051	0.0611

B.5.6 Each of the independent components of X follows a shifted gamma distribution. It now remains to find the appropriate α parameters for each component. My first attempt was to calibrate the model to the skewness of the transformed series. Unfortunately the log return on gilts had a negative skewness which is inconsistent with the assumptions that the $\gamma_{ij} \geq 0$. I therefore abandoned this idea, and instead considered the kurtosis. If I can fit this, then at least the sizes of the jumps are realistic, even if the directions do not fit perfectly. This worked fine for the first two components, but when examining the second two, corresponding to the regression residuals, I find a negative kurtosis. I suspect that this is a noise effect arising from the small sample size, since negative kurtosis is inconsistent with random walks in general, not just with the Γ process. In these cases I tried to make the kurtosis as small as possible, which means large values of α. In fact, for large α the model tends towards a structure based on the multivariate normal distribution. I took $\alpha=10$, which is large enough to be indistinguishable from the normal for most practical purposes.

B.5.7 Readers will notice that the above methods are rather *ad hoc*, and difficult to justify from a statistical perspective. I had to resort to these because the more conventional techniques did not work. I have also failed to derive standard errors and significance tests. In one sense, these would be useless anyway, because, in order to derive the test, I would have to assume the null hypothesis based on independent gamma distributions. I do not seriously believe that this hypothesis could be true. The really interesting question is whether a model fitted using my methods is useful for actuarial purposes when the underlying world follows some much more complex rules. I have no doubt that my statistical methods can be improved, and I look forward to hearing suggested improvements from readers of this paper.

B.6 *Estimation of Mean Terms*

B.6.1 I have applied economic theory rather than historical data to determine the mean returns for each asset class. The simplest approach is that followed by Dyson & Exley, using the rational expectation approach. If these estimates are unbiased, then there is no explicit risk premium loaded into any asset class. All asset classes have the same mean log return.

B.6.2 For the random walk model, and for my own model, I have constructed means via an equilibrium argument. The idea is to specify a particular class of investor, with defined utility function, and also to specify the portfolio that investor optimally holds. The assumptions I have made, somewhat arbitrarily, are as follows:

	Portfolio	Proportions	Optimal	Durations
Gilts	p_1	15%	t_1	15
Equities	p_2	60%	t_2	0
Property	p_3	20%	t_3	0
Index-linked	p_4	5%	t_4	10

I assume that this investor has a power law utility as in ¶A.5.5 with κ =3, and with all returns measured relative to the return on cash. Various empirical studies on the appropriate value of κ exist in the literature, and the answers are not totally consistent or conclusive. A good reference on this subject is Mehra & Prescott (1985). The weights for the different time horizons turn out not to affect the calibration, so I make no particular assumptions.

B.6.3 Under these assumptions, and following the method of Appendix A for the optimisation, I obtain the following mean real returns for the random walk model:

Equities	9.89%
Gilts	4.50%
Cash	1.45%
Property	5.67%
Index-linked	4.00%

For my own jump equilibrium model I obtain the following parameter estimates:

		α		
	4.9284	10.0000	6.7772	10.0000
		β		
Sterling	0.0623	0.1185	0.0326	0.0029
RPI	0.0623	0.1185	0.0326	0.0029
Equities	0.0410	0.0596	0.0326	0.0029
Property	0.0630	0.1048	−0.0212	0.0029
		γ		
Sterling	0.0033	0.0005	0.0000	0.0000
RPI	0.0016	0.0004	0.0002	0.0019
Equities	0.0000	0.0000	0.0000	0.0000
Property	0.0000	0.0000	0.0000	0.0000

B.6.4 It is of interest to consider long run returns relative to cash for each of these models. Financial economists would usually look at arithmetic average returns. However, the actuarial tradition is to look at geometric returns, so that is what I have reluctantly done here. The results I obtain (expressed as an excess annual return relative to cash) are:

	Historic	Random walk	Dyson & Exley	Jump equilibrium	Wilkie
Gilts	0.22%	3.09%	0.00%	2.02%	1.38%
Equities	6.02%	8.80%	0.00%	6.56%	3.94%
Property	−3.35%	4.31%	0.00%	1.94%	2.28%
Index-linked	1.98%	2.58%	0.00%	0.88%	2.23%

I have derived these analytically, except for the Wilkie model, where I used simulations. The chaotic and fractal models will give the same result as the random walk model.

B.6.5 I have considered a particular class of investor and estimated the objective function that the investor is optimising. I then equated the actual portfolio held by this benchmark investor to the theoretical optimum under a stochastic model. This, too, is subject to significant estimation difficulties, particularly regarding the utility function. However, when I come to optimise on behalf of my client, there will also be uncertainties in my client's utility function. The answers to the optimisation are sensitive, not so much to the utility in absolute terms, but to the difference between the assumed utilities of the client and the benchmark investor. I believe that I can estimate these relativities with some degree of confidence. Thus, if, for example, my client is much more concerned about adversely high inflation than would be typical in the industry, I would recommend my client to be disproportionately weighted towards index-linked gilts in his asset portfolio, relative to his peers.

APPENDIX C

COMPUTER CODE

C.1 *Structure of the Programs*

C.1.1 This appendix contains all the code that I have used to generate all the examples in Section 3. This includes the code for running the stochastic models, determining efficient frontiers and dynamic optimisation. I believe that the only way to start to understand many of these models is by actually running them, and that is the spirit in which I have made my code public.

C.1.2 I have programmed the models in Microsoft Excel version 5 for Windows. This version of Excel comes with a version of the Visual Basic programming language which links seamlessly with Excel. In particular, calls to Visual Basic functions can return values directly to a spreadsheet. This appendix contains the Visual Basic functions, but not the spreadsheets.

C.1.3 The reasons for choosing Excel were several. Excel is a popular and inexpensive package which can run under several hardware and operating system platforms. The Visual Basic code is familiar to many, and is reasonably easy for a novice to learn and understand. The interface with Excel makes preparation of graphs and charts very easy. I do not claim that my code is particularly efficient, and sometimes I have deliberately avoided programming short cuts in order to achieve greater readability. I could have achieved a faster run time if I had used a compiled language such as C or FORTRAN, but I find these languages less readable and harder to debug. In any case, the generation of stochastic scenarios is not generally a significant delay in real projects; the time-consuming part is feeding the scenarios through a model office or equivalent, so speed improvements in scenario generation are of little help in the overall scheme of things.

C.1.4 Visual Basic has a number of syntactical features whose meaning may not be obvious at first sight. The 'Type' construction defines a data structure (similar to a 'struct' in C). The default when calling functions is to pass arguments by reference, so that if their values are changed within a function, the new values are written back to the calling routine. I have chosen to declare all variables using the 'Dim' statement; this is not necessary for the code to run, but makes it easier to debug. Dynamic memory space is allocated using 'ReDim'. Variables declared with 'Static' retain their values between function calls, exactly as with the corresponding syntax in C. The 'Variant' type is a built-in data type which can store any other built-in type, including arrays of built-in types, and also remembers what type it is.

C.1.5 I have organised the code into a series of *modules*. Each module contains conceptually related functions, but the grouping has little computational significance. However, constants can be defined at a module level; such constants are accessible to functions and subroutines in that module, but not elsewhere. I have described the modules in a 'bottom up' style, starting with the simplest functions and moving on to more complex operations that call the simpler ones. Figure C.1.5 shows the main dependencies between principal functions.

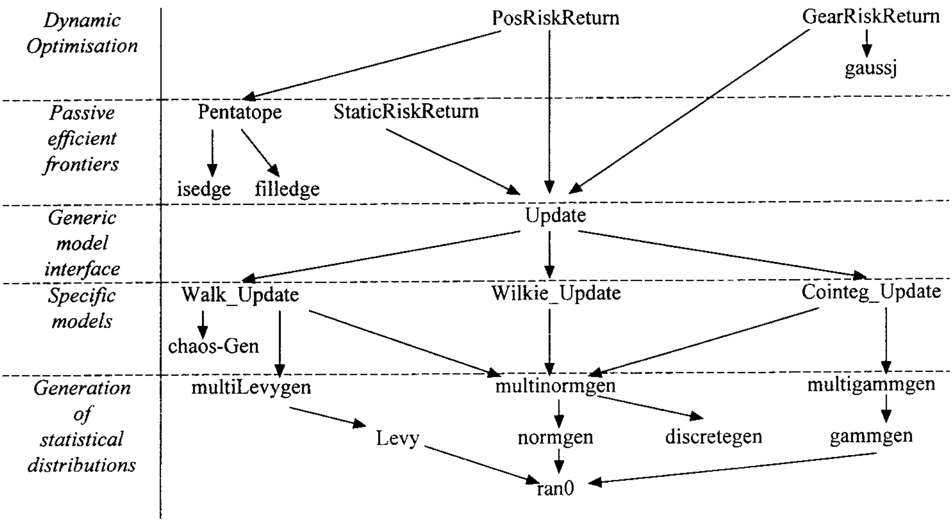

Figure C.1.5. Functional dependencies

In this diagram, an arrow links function calls, so that 'foo → bar' means that the function 'foo' calls the function 'bar'.

C.2 *Generation of Statistical Distributions*

C.2.1 I have grouped together the functions which generate random variables from the distributions required for the models. These include normal, gamma and Lévy stable distributions. To these I have added simple routines to produce arrays of such distributions. The normal distribution is adapted from Flannery *et al.* (1989), while the Lévy distribution is taken from Chambers *et al.* (1973). The gamma distribution algorithm is my own.

C.2.2 I have also coded some discrete approximations to the normal distribution. I use these for carrying out numerical integration in the dynamic optimisation.

C.2.3 Finally, I have included a routine to produce uniform random variables. This is based on the system supplied routine, which comes from a generic type known as a *linear congruential generator.* This type of generator suffers from certain autocorrelation problems; my algorithm, which is adapted from Flannery *et al.* (1989) simply scrambles the order of the points supplied by the system and eliminates some of its difficulties.

C.2.4 The code for statistical distributions follows:

```
Type multi_error
   Z() As Double
End Type

Function normgen() As Double
'Return sample from N(0,1) distribution
Static got_one As Boolean, stored As Double
Dim x1 As Double, x2 As Double, r As Double
If got_one Then
   normgen = stored
   got_one = False
Else
   'generate a point in the unit square
   Do
      x1 = 2 * Ran0 - 1
      x2 = 2 * Ran0 - 1
      r = x1 ^ 2 + x2 ^ 2
   'reject unless inside unit circle
   Loop While r >= 1 Or r = 0
   r = Sqr(-2 * Log(r) / r)
   stored = x1 * r
   normgen = x2 * r
   got_one = True
End If
End Function

Sub multinormgen(Z As multi_error, Optional discrete)
Dim i As Integer
Static beenherebefore As Boolean, H() As multi_error

If IsMissing(discrete) Then
   'random generation
   For i = 1 To UBound(Z.Z)
      Z.Z(i) = normgen
   Next i
Else
   If beenherebefore Then
      'no need to initialise H
   Else
      ReDim H(1 To UBound(Z.Z) + 2)
```

```
      discretegen H()
      beenherebefore = True
   End If
   Z = H(discrete)
End If
End Sub

Function gammgen(alpha) As Double
'Return sample from Gamma(alpha) distribution
Dim beta As Double, gamma As Double
Dim u1 As Double, u2 As Double, G As Double
beta = Sqr(alpha + 0.25) - 0.5
gamma = beta + 1
Do
   u1 = Ran0
   u2 = Ran0
   G = u1 ^ (1 / beta) / u2 ^ (1 / gamma) / 2.71828 * _
      (alpha + gamma) ^ ((alpha + gamma) / _
      (beta + gamma)) / (alpha - beta) _
      ^ ((alpha - beta) / (beta + gamma))
Loop Until u1 * u2 <= (2.71828 / (alpha - beta)) ^ _
   (gamma * (alpha - beta) / (beta + gamma)) * _
   (2.71828 / (alpha + gamma)) ^ (beta * (alpha + _
   gamma) / (beta + gamma)) * G ^ alpha * Exp(-G)
gammgen = G
End Function

Sub multigammgen(Z As multi_error, alpha)
Dim i As Integer
For i = 1 To UBound(Z.Z)
   Z.Z(i) = gammgen(alpha(i))
Next i
End Sub

Function Levy(alpha, beta) As Double
'return sample from Levy stable (alpha, beta)
Dim phi0 As Double, phi As Double, W As Double
phi0 = (alpha - 2) / 2 / alpha * 3.1415927 * beta
phi = 3.1415927 * (Ran0 - 0.5)
W = -Log(Ran0)
Levy = Sin(alpha * (phi - phi0)) / Cos(phi) ^ _
   (1 / alpha) * (Cos(phi - alpha * (phi - phi0)) / W) ^ _
   ((1 - alpha) / alpha)
End Function
```

```
Sub multiLevygen(Z As multi_error, alpha, beta)
Dim i As Integer
For i = 1 To UBound(Z.Z)
   Z.Z(i) = Levy(alpha, beta)
Next i
End Sub

Function discretegen(Z() As multi_error)
'generates a discrete approximationn to the normal
Dim p As Integer, i As Integer, j As Integer
p = UBound(Z) \ 2 - 1
For j = 1 To 2 * p + 2
   ReDim Z(j).Z(1 To 2 * p)
   With Z(j)
      For i = 1 To p
         .Z(i) = 1.414214 * Sin(i * j * 3.141593 / (p + 1))
         .Z(i + p) = 1.414214 * Cos(i * j * 3.141593 / _
            (p + 1))
      Next i
   End With
Next j
End Function

Function Ran0(Optional seed) As Double
   'returns a uniform (0,1) random variable
   'obtained by scrambling system routine.
Static v(1 To 20) As Double, beenherebefore As Boolean
Static y As Double
Dim dummy As Double, j As Integer
If IsMissing(seed) Then
   If beenherebefore Then
      'no need to initialise
   Else
      dummy = Rnd(-1)
      For j = 1 To 20
         v(j) = Rnd()
      Next j
         y = Rnd()
   End If
Else
   'reinitialise with seed provided
   dummy = Rnd(-1.414 - 1.723 * Abs(seed) ^ 2.718)
   dummy = Rnd()
   dummy = Rnd()
```

```
   For j = 1 To 20
      v(j) = Rnd()
   Next j
   y = Rnd()
End If
beenherebefore = True
j = 1 + Int(20 * y)
y = v(j)
Ran0 = y
v(j) = Rnd()
End Function
```

C.3 *Specific Models*

C.3.1 Inevitably, there is some code which is model specific. Surprisingly, this is a relatively small part of the whole, as many of the models have common features which means that much of the code can be re-used. I have divided the models into three types: those based on the random walk (including the chaotic and fractal models), the Wilkie model, and lastly the cointegrated models.

C.3.2 The code for the random walk and associated models follows. The main differences between the models arise from the error terms:

```
Const QMU = 0.047
Const QA = 0.58

Sub Chaos_gen(Z As multi_error)
Static beenherebefore As Boolean
Dim i As Integer
If beenherebefore Then
   For i = 1 To UBound(Z.Z)
      Z.Z(i) = 1.414 * (1 - Z.Z(i) ^ 2)
   Next i
Else
   For i = 1 To UBound(Z.Z)
      Z.Z(i) = 2.818 * Ran0 - 1.414
   Next i
   beenherebefore = True
End If
End Sub

Sub Walk_Update(S As Return_Index, Model As String, _
   Optional discrete)
Static inflation_rate As Double, mu(1 To 6) As Double
Static C(1 To 6, 1 To 6) As Double, Z As multi_error
```

```
Static initialised As Boolean
Dim oldZ As multi_error
Dim i As Integer, j As Integer, E As multi_error
Dim oldinflation_rate As Double
If Not initialised Then
   mu(1) = 0
   mu(2) = 0.0989
   mu(3) = 0.045
   mu(4) = 0.0145
   mu(5) = 0.0567
   mu(6) = 0.04
   C(1, 1) = 0.056
   C(1, 2) = 0
   C(1, 3) = 0
   C(1, 4) = 0
   C(1, 5) = 0
   C(1, 6) = 0
   C(2, 1) = -0.0529
   C(2, 2) = 0.2196
   C(2, 3) = 0
   C(2, 4) = 0
   C(2, 5) = 0
   C(2, 6) = 0
   C(3, 1) = -0.0716
   C(3, 2) = 0.0702
   C(3, 3) = 0.1045
   C(3, 4) = 0
   C(3, 5) = 0
   C(3, 6) = 0
   C(4, 1) = -0.0406
   C(4, 2) = 0.0016
   C(4, 3) = 0.0058
   C(4, 4) = 0.0332
   C(4, 5) = 0
   C(4, 6) = 0
   C(5, 1) = -0.0488
   C(5, 2) = 0.0743
   C(5, 3) = -0.0676
   C(5, 4) = -0.0619
   C(5, 5) = 0.2056
   C(5, 6) = 0
   C(6, 1) = -0.0164
   C(6, 2) = 0.0617
   C(6, 3) = 0.0952
```

```
   C(6, 4) = -0.0006
   C(6, 5) = 0.0049
   C(6, 6) = 0.0558
   ReDim Z.Z(1 To 6)
   Chaos_gen Z
   initialised = True
End If
If S.Year = 0 Then
   inflation_rate = QMU
End If
ReDim E.Z(1 To 6)
oldinflation_rate = inflation_rate
oldZ = Z
If Model = "Chaotic" Then
   Chaos_gen Z
ElseIf Model = "Fractal" And IsMissing(discrete) Then
   multiLevygen Z, 1.65, -0.11
Else
   multinormgen Z, discrete
End If
For i = 1 To 6
   E.Z(i) = 0
   For j = 1 To i
      E.Z(i) = E.Z(i) + C(i, j) * Z.Z(j)
   Next j
Next i
inflation_rate = QA * inflation_rate + (1 - QA) * QMU + E.Z(1)

With S
   .Retail_Price = .Retail_Price * Exp(inflation_rate)
   .Equity = .Equity * Exp(mu(2) + E.Z(2) + _
      inflation_rate)
   .Fixed_Interest = .Fixed_Interest * Exp(mu(3) _
      + E.Z(3) + inflation_rate)
   .Cash = .Cash * Exp(mu(4) + E.Z(4) + inflation_rate)
   .Property = .Property * Exp(mu(5) _
      + E.Z(5) + inflation_rate)
   .Index_Linked = .Index_Linked * Exp(mu(6) _
      + E.Z(6) + inflation_rate)
   .Real_Yield = 0.04
End With
If IsMissing(discrete) Then
   'leave updated inflation rate
Else
```

```
   'restore previous state variables
   inflation_rate = oldinflation_rate
   Z = oldZ
End If
End Sub
```

C.3.3 We now move on to the particular code for the Wilkie model. Paragraph numbers refer to Wilkie (1995):

```
Type Wilkie_State
'contains the stationary state variables
'starting with the AR1 variables
   i As Double
   YN As Double
   CN As Double
   BD As Double
   lnZ As Double
   lnR As Double
'now the moving averages of inflation
   DM As Double
   CM As Double
   EM As Double
'and finally, the error terms carried forward
   YE As Double
   DE As Double
End Type

'retail price inflation parameters section 2.3.7
Const QMU As Double = 0.047
Const QA As Double = 0.58
Const QSD As Double = 0.0425

'dividend yields, section 4.3.4
Const YW As Double = 1.8
Const YA As Double = 0.55
'Const YMU As Double = 0.0375 this is always logged,
'so store logged value
Const lnYMU As Double = -3.283414346
Const YSD As Double = 0.155

'dividends, section 5.3.8
Const DW As Double = 0.58
Const DD As Double = 0.13
Const DMU As Double = 0.016
```

```
Const DY As Double = -0.175
Const DB As Double = 0.57
Const DSD As Double = 0.07

'long term interest rates, section 6.3.7
Const CD As Double = 0.045
Const CMU As Double = 0.0305
Const CA1 As Double = 0.9
Const CY As Double = 0.34
Const CSD As Double = 0.185

'short term interest rates, section 7.1.7
Const BMU As Double = 0.23
Const BA As Double = 0.74
Const BSD As Double = 0.18

'property yields, section 8.2.4
'Const ZMU As Double = 0.074 this is always logged,
' so use log value instead
'Const lnZMU As Double = -2.603690186
'Const ZA As Double = 0.91
'Const ZSD As Double = 0.12
'Replace by Daykin & Hey figures
Const lnZMU As Double = -2.995732274
Const ZA As Double = 0.6
Const ZSD As Double = 0.075

'property rents, section 8.3.3
'Daykin & Hey figures for means
'but retain Wilkie for variance
Const EW As Double = 1
Const ED As Double = 0.1 ' Wilkie suggests 0.13 instead
Const EMU As Double = -0.01 ' Wilkie suggests 0.003
Const EBZ As Double = 0.24
Const ESD As Double = 0.06

'index linked gilt yields, section 9.2.4
'Const RMU As Double = 0.04 always logged so store log
Const lnRMU As Double = -3.218875825
Const RA As Double = 0.55
Const RBC As Double = 0.22
Const RSD As Double = 0.05

'initialisation of states
```

```
Sub Neutral(W As Wilkie_State)
With W
'starting with the AR1 variables
   .i = QMU
   .YN = lnYMU
   .CN = 0
   .BD = BMU
   .lnZ = lnZMU
   .lnR = lnRMU
'now the moving averages of inflation
   .DM = QMU
   .CM = QMU
   .EM = QMU
'and finally, the error terms carried forward
   .YE = 0
   .DE = 0
End With
End Sub

Sub Preindex(S As Return_Index, W As Wilkie_State)
Dim C As Double
With W
   S.Equity = S.Equity * Exp(YW * .i + .YN _
      + DY * .YE + DB * .DE)
   C = .CM + CMU * Exp(.CN)
   S.Fixed_Interest = S.Fixed_Interest * C
   S.Cash = S.Cash * (1 + C * Exp(-.BD))
   S.Index_Linked = S.Index_Linked * Exp(.lnR)
   S.Property = S.Property * Exp(.lnZ)
End With
End Sub

Sub Advance(W As Wilkie_State, Z As multi_error)
With W
   .i = QMU + QA * (.i - QMU) + QSD * Z.Z(1)
   .YE = YSD * Z.Z(2)
   .YN = lnYMU + YA * (.YN - lnYMU) + .YE
   .DM = DD * .i + (1 - DD) * .DM
   .DE = DSD * Z.Z(3)
   .CM = CD * .i + (1 - CD) * .CM
   .CN = CA1 * .CN + CY * .YE + CSD * Z.Z(4)
   .BD = BMU + BA * (.BD - BMU) + BSD * Z.Z(5)
   .lnZ = lnZMU + ZA * (.lnZ - lnZMU) + ZSD * Z.Z(6)
   .EM = ED * .i + (1 - ED) * .EM
```

```
   .lnR = lnRMU + RA * (.lnR - lnRMU) _
      + RBC * CSD * Z.Z(4) + RSD * Z.Z(8)
End With
End Sub

Sub Postindex(S As Return_Index, W As Wilkie_State, _
   Z As multi_error)
With W
   S.Retail_Price = S.Retail_Price * Exp(.i)
   S.Equity = S.Equity * Exp(DW * .DM + (1 - DW) * .i _
      + DMU + .DE) * (Exp(-YW * .i - .YN) + 1)
   S.Fixed_Interest = S.Fixed_Interest _
      * (1 / (.CM + CMU * Exp(.CN)) + 1)
   S.Index_Linked = S.Index_Linked * Exp(.i) _
      * (Exp(-.lnR) + 1)
   S.Real_Yield = Exp(.lnR)
   S.Property = S.Property * Exp(EW * .EM + (1 - EW) * _
      .i + EMU + EBZ * ZSD * Z.Z(6) + ESD * Z.Z(7)) * _
      (Exp(-.lnZ) + 1)
End With
End Sub

Sub Wilkie_Update(S As Return_Index, Optional discrete)
'if using discrete approximation then do not update W
Static W As Wilkie_State
Dim Z As multi_error, Wtemp As Wilkie_State
ReDim Z.Z(1 To 8)
If S.Year = 0 Then
   Neutral W
End If
Wtemp = W
Preindex S, W
multinormgen Z, discrete
Advance W, Z
Postindex S, W, Z
If IsMissing(discrete) Then
   'leave W updated
Else
   'restore original state variable
   W = Wtemp
End If
End Sub
```

C.3.4 The last of the specific model classes consists of the cointegrated models. Once again I have attempted to re-use as much of the code as possible. This means that the similarity of the models is clear, but the code for the Dyson & Exley model is rather more longwinded than is absolutely necessary:

```
Const STDTERM = 3

Type Cointeg_State
'contains the state variables
      G(1 To 4) As Double
      H(1 To 4) As Double
End Type

Sub Cointeg_Update(S As Return_Index, Model As String, _
   Optional discrete)

Static alpha(1 To 4) As Double
Static beta(1 To 4, 1 To 4) As Double
Static gamma(1 To 4, 1 To 4) As Double
Static x As Cointeg_State
Static init_yield(1 To 4) As Double
Static duration(1 To 4) As Double
Dim Z As multi_error, Xtemp As Cointeg_State
Dim i As Integer, j As Integer, t As Integer
Dim price(1 To 4), rollup(1 To 4), bond(1 To 4)
Dim yield(1 To 4)
'price relates to the spot price of an asset
'rollup is a total return index of zero duration bonds
'bond is total return of a constant maturity index.
'yield is the yield on a standard term bond

If S.Year = 0 Then
   'initialise everything
   alpha(1) = 4.9284
   alpha(2) = 10
   alpha(3) = 6.7772
   alpha(4) = 10
   beta(1, 1) = 0.0623
   beta(1, 2) = 0.1185
   beta(1, 3) = 0.0326
   beta(1, 4) = 0.0029
   beta(2, 1) = 0.0623
   beta(2, 2) = 0.1185
   beta(2, 3) = 0.0326
```

```
   beta(2, 4) = 0.0029
   beta(3, 1) = 0.041
   beta(3, 2) = 0.0596
   beta(3, 3) = 0.0326
   beta(3, 4) = 0.0029
   beta(4, 1) = 0.063
   beta(4, 2) = 0.1048
   beta(4, 3) = -0.0212
   beta(4, 4) = 0.0029
   gamma(1, 1) = 0.0033
   gamma(1, 2) = 0.0005
   gamma(1, 3) = 0
   gamma(1, 4) = 0
   gamma(2, 1) = 0.0016
   gamma(2, 2) = 0.0004
   gamma(2, 3) = 0.0002
   gamma(2, 4) = 0.0019
   gamma(3, 1) = 0
   gamma(3, 2) = 0
   gamma(3, 3) = 0
   gamma(3, 4) = 0
   gamma(4, 1) = 0
   gamma(4, 2) = 0
   gamma(4, 3) = 0
   gamma(4, 4) = 0
   init_yield(1) = 2 * Log(1 + 0.075 / 2)
   init_yield(2) = 2 * Log(1 + 0.035 / 2)
   init_yield(3) = 0
   init_yield(4) = 0
   duration(1) = 15
   duration(2) = 10
   duration(3) = 0
   duration(4) = 0
   With x
      For i = 1 To 4
         .G(i) = 0
         .H(i) = 0
      Next i
   End With
End If
Xtemp = x
ReDim Z.Z(1 To 4) As Double
If Model = "JumpEq" And IsMissing(discrete) Then
   multigammgen Z, alpha
```

```
Else
   multinormgen Z, discrete
   With Z
      For i = 1 To 4
         'normal aproximation to the gamma
         .Z(i) = alpha(i) + Sqr(alpha(i)) * .Z(i)
      Next i
   End With
End If
For i = 1 To 4
   'update gamma processes and integrated processes
   ' using trapezium rule
   With x
      .H(i) = .H(i) + .G(i) / 2
      .G(i) = .G(i) + Z.Z(i)
      .H(i) = .H(i) + .G(i) / 2
   End With
Next i
For i = 1 To 4
   'calculate time. Note that working to project next year
   t = S.Year + 1
   'all these quantities are stored as logs
   price(i) = -t * init_yield(i)
   'price of asset exluding investment return
   yield(i) = init_yield(i)
   rollup(i) = 0 'price of rollup with zero duration
   bond(i) = 0 'price of rollup with specified duration
   For j = 1 To 4
      If Model = "JumpEq" Then
         If gamma(i, j) = 0 Then
            price(i) = price(i) + alpha(j) _
               * t * Log(1 - beta(i, j))
         Else
            price(i) = price(i) - alpha(j) * t + alpha(j) _
               / gamma(i, j) * ((1 - beta(i, j) + _
               gamma(i, j) * t) * Log(1 - beta(i, j) + _
               gamma(i, j) * t) - (1 - beta(i, j)) * _
               Log(1 - beta(i, j)))
            yield(i) = yield(i) + alpha(j) / gamma(i, j) _
               / STDTERM * ((1 - beta(i, j) + _
               gamma(i, j) * t) * Log(1 - beta(i, j) + _
               gamma(i, j) * t) - (1 - beta(i, j)) * _
               Log(1 - beta(i, j)) + (1 - beta(i, j) + _
               gamma(i, j) * STDTERM) * Log(1 - beta(i, j) + _
```

```
                gamma(i, j) * STDTERM) - (1 - beta(i, j) + _
                gamma(i, j) * (t + STDTERM)) * _
                Log(1 - beta(i, j) + gamma(i, j) * _
                (t + STDTERM)))
        End If
        rollup(i) = rollup(i) + alpha(j) * t * _
           Log(1 - beta(i, j))
        bond(i) = bond(i) + alpha(j) * t * _
           Log(1 - beta(i, j) + duration(i) * gamma(i, j))
     ElseIf Model = "Dyson&Exley" Then
        'first order approximation
        'for small beta and gamma
        price(i) = price(i) - alpha(j) * (beta(i, j) _
           * t - gamma(i, j) * t ^ 2 / 2)
        rollup(i) = rollup(i) - alpha(j) * _
           beta(i, j) * t
        bond(i) = bond(i) - alpha(j) * t * _
           (beta(i, j) - duration(i) * gamma(i, j))
        yield(i) = yield(i) - gamma(i, j) * _
           alpha(j) * t
     End If
     'now apply stochastic terms
     With x
        price(i) = price(i) + beta(i, j) * .G(j) _
           - gamma(i, j) * .H(j)
        rollup(i) = rollup(i) + beta(i, j) * .G(j)
        bond(i) = bond(i) + (beta(i, j) - _
           duration(i) * gamma(i, j)) * .G(j)
        yield(i) = yield(i) + gamma(i, j) * .G(j)
     End With
  Next j
Next i
'Now poke answer back into S
With S
   .Retail_Price = Exp(price(2) - price(1))
   .Equity = Exp(rollup(3) - price(1))
   .Fixed_Interest = Exp(bond(1) - price(1))
   .Cash = Exp(rollup(1) - price(1))
   .Index_Linked = Exp(bond(2) - price(1))
   .Real_Yield = Exp(yield(2)) - 1
   .Property = Exp(rollup(4) - price(1))
End With
If IsMissing(discrete) Then
   'leave X updated
```

```
Else
   'restore original state variable
   x = Xtemp
End If
End Sub
```

C.4 *Generic Model Interface*

C.4.1 I have collected together the features common to all models in a single module. This defines a number of data structures, and the simple generic code for calculating investment indices year by year. A function is also provided which puts sample paths into a spreadsheet array. One slightly subtle twist is the use of the optional parameter `discrete`. If this parameter is omitted, then the code generates a random scenario, and the state variables are updated to prepare for the following year. On the other hand, if `discrete` is provided, then a carefully selected sample outcome is returned, without the state variables being updated. For a dynamic optimisation, the procedures are called several times with `discrete` set in order to plan a suitable matching allocation for the following year, and then finally the actual outcome is simulated using a call without `discrete`.

C.4.2 The code follows below:

```
Type Return_Index
   'These are all total return indices
   'Including income gross of tax
   'with the exception of RPI, which has no income
   'and real yield, which is for calculating reserves
   Year As Integer
   Retail_Price As Double
   Equity As Double
   Fixed_Interest As Double
   Cash As Double
   Index_Linked As Double
   Property As Double
   Real_Yield As Double
End Type

Sub Unit(S As Return_Index)
   With S
      .Year = 0
      .Retail_Price = 1
      .Equity = 1
      .Fixed_Interest = 1
      .Cash = 1
      .Index_Linked = 1
```

```
      .Property = 1
      .Real_Yield = 0.04
   End With
End Sub

Sub Fill(index_array As Variant, S As Return_Index)
   With S
      index_array(.Year, 1) = .Year
      index_array(.Year, 2) = .Retail_Price
      index_array(.Year, 3) = .Equity
      index_array(.Year, 4) = .Fixed_Interest
      index_array(.Year, 5) = .Cash
      index_array(.Year, 6) = .Index_Linked
      index_array(.Year, 7) = .Property
   End With
End Sub

Sub Update(S As Return_Index, Model As String, Optional discrete)
   'advances stochastic model by 1 year.
   'if discrete is defined, then this is just a dummy run,
   'to test one of a few scenarios. State variables are
   'preserved (but S is updated). If discrete is missing
   'then this is a real life simulation and state
   'variables are updated.
   Select Case Model
   Case "RandomWalk", "Fractal", "Chaotic"
      Walk_Update S, Model, discrete
      'fortunately, if discrete was missing when Update
      'was called, VBA remembers this, and Walk_update
      'also knows it is missing.
   Case "Wilkie"
      Wilkie_Update S, discrete
   Case "JumpEq", "Dyson&Exley"
      Cointeg_Update S, Model, discrete
   End Select
   S.Year = S.Year + 1
End Sub

Function Scenario(Model As String, horizon As Integer, _
   seed As Double) As Variant
   'Skeleton function for putting model output
   'into an array. Use seed for random number generator
   'with arbitrary non-round coefficients
   'Set out range for output
```

```
   ReDim outscen(0 To horizon, 1 To 7) As Double
   Dim S As Return_Index, dummy As Double
   dummy = Ran0(seed)
   Unit S
   Fill outscen, S
   Do
      Update S, Model
      Fill outscen, S
   Loop Until S.Year = horizon
   Scenario = outscen
End Function
```

C.5 *Passive Efficient Frontiers*

C.5.1 The next section contains the code for plotting efficient frontiers. This includes the selection of asset mixes to be tried, as well as the code which rolls up investment performance to determine simulated outcomes, posting back standard deviations and means:

```
Type Asset_Mix
   'composition of a portfolio
   Equity As Double
   Fixed_Interest As Double
   Cash As Double
   Index_Linked As Double
   Property As Double
End Type

Function StaticRiskReturn(Model As String, _
   horizon As Integer, seed As Double, _
   alloc_range As Variant, nosims As Integer, _
   Optional cap) As Variant
'calculates mean and standard deviation of return
'for a range of static investment strategies.
Dim nostrats As Integer, strat As Integer, t As Integer
nostrats = alloc_range.Rows.Count
Dim S As Return_Index, oldS As Return_Index, sim As Integer,
dummy As Double
Dim i As Integer, j As Integer
ReDim alloc_array(1 To nostrats) As Asset_Mix
ReDim asset_share(1 To nostrats) As Double
ReDim sumreturn(1 To nostrats) As Double
ReDim outvec(1 To nostrats, 1 To 2) As Double
ReDim sumsquare(1 To nostrats) As Double
'Read allocations into alloc_array
```

```
dummy = Ran0(seed)
For i = 1 To nostrats
   With alloc_array(i)
      .Equity = alloc_range(i, 1).Value
      .Fixed_Interest = alloc_range(i, 2).Value
      .Cash = alloc_range(i, 3).Value
      .Index_Linked = alloc_range(i, 4).Value
      .Property = alloc_range(i, 5).Value
   End With
Next i
For sim = 1 To nosims
   Unit S
   For i = 1 To nostrats
      asset_share(i) = 0
   Next i
   Do
      oldS = S
      Update S, Model
   For i = 1 To nostrats
      asset_share(i) = asset_share(i) + oldS.Retail_Price
      With alloc_array(i)
         asset_share(i) = asset_share(i) * ( _
            .Equity * S.Equity / oldS.Equity _
            + .Fixed_Interest * S.Fixed_Interest _
            / oldS.Fixed_Interest _
            + .Cash * S.Cash / oldS.Cash _
            + .Index_Linked * S.Index_Linked _
            / oldS.Index_Linked _
            + .Property * S.Property / oldS.Property)
      End With
      Next i
   Loop Until S.Year = horizon
   For i = 1 To nostrats
      If IsMissing(cap) Then
         'no maximum applies
      Else
         If asset_share(i) > cap(i) * S.Retail_Price Then
            asset_share(i) = cap(i) * S.Retail_Price
         End If
      End If
      sumreturn(i) = sumreturn(i) + asset_share(i) / _
         S.Retail_Price
      sumsquare(i) = sumsquare(i) + (asset_share(i) / _
         S.Retail_Price) ^ 2
```

```
   Next i
Next sim
For i = 1 To nostrats
   outvec(i, 1) = Sqr((sumsquare(i) - sumreturn(i) ^ 2 _
      / nosims) / (nosims - 1))
   outvec(i, 2) = sumreturn(i) / nosims
Next i
StaticRiskReturn = outvec

End Function

Function pentatope(partitions As Integer) As Variant
'produces asset allocations when the portfolio is split into n
partitions.
Dim comb As Integer, nocombs As Integer
nocombs = (partitions + 1) * (partitions + 2) * _
   (partitions + 3) * (partitions + 4) / 24 + 6
ReDim alloc_range(1 To nocombs, 1 To 5) As Double
Dim i1 As Integer, i2 As Integer, i3 As Integer
Dim i4 As Integer
comb = 1
fill_edge alloc_range, comb, partitions, 1, 2
fill_edge alloc_range, comb, partitions, 2, 3
fill_edge alloc_range, comb, partitions, 3, 4
fill_edge alloc_range, comb, partitions, 4, 5
fill_edge alloc_range, comb, partitions, 5, 1
fill_edge alloc_range, comb, partitions, 1, 3
fill_edge alloc_range, comb, partitions, 3, 5
fill_edge alloc_range, comb, partitions, 5, 2
fill_edge alloc_range, comb, partitions, 2, 4
fill_edge alloc_range, comb, partitions, 4, 1
For i4 = 0 To partitions
   For i3 = 0 To i4
      For i2 = 0 To i3
         For i1 = 0 To i2
            If isedge(i1, i2, i3, i4, partitions) Then
               'skip. We have already covered this
            Else
               comb = comb + 1
               alloc_range(comb, 1) = i1 / partitions
               alloc_range(comb, 2) = (i2 - i1) / partitions
               alloc_range(comb, 3) = (i3 - i2) / partitions
               alloc_range(comb, 4) = (i4 - i3) / partitions
               alloc_range(comb, 5) = 1 - i4 / partitions
```

```
            End If
         Next i1
      Next i2
   Next i3
Next i4
pentatope = alloc_range
End Function

Function isedge(i1 As Integer, i2 As Integer, _
   i3 As Integer, i4 As Integer, partitions) As Boolean
'this is an edge if at least three of the asset classes 'are
absent
Dim no_absent As Integer
If i1 = 0 Then no_absent = no_absent + 1
If i1 = i2 Then no_absent = no_absent + 1
If i2 = i3 Then no_absent = no_absent + 1
If i3 = i4 Then no_absent = no_absent + 1
If i4 = partitions Then no_absent = no_absent + 1
If no_absent >= 3 Then
   isedge = True
Else
   isedge = False
End If
End Function

Sub fill_edge(alloc_range, comb As Integer, _
   partitions As Integer, start As Integer, _
   finish As Integer)
Dim i As Integer, j As Integer
For i = 0 To partitions
   If i > 0 Then
      comb = comb + 1
   End If
   For j = 1 To 5
      If j = start Then
         alloc_range(comb, j) = 1 - i / partitions
      ElseIf j = finish Then
         alloc_range(comb, j) = i / partitions
      Else
         alloc_range(comb, j) = 0
      End If
   Next j
Next i
End Sub
```

C.6 *Dynamic Optimisation Functions*

C.6.1 Last, but not least, we come to the algorithms for dynamic optimisation. This contains two algorithms. The first examines positive allocation by working through a list of possible allocations and choosing the best from each list. The second algorithm allows for short positions and gearing, and solves for the optimal portfolio over a single-year horizon by Gaussian elimination. The Gaussian elimination routine is adapted from Flannery *et al.* (1989).

C.6.2 The code follows:

```
Function PosRiskReturn(Model As String, _
   horizon As Integer, seed As Double, _
   nosims As Integer, target_range As Variant) _
   As Variant
'crude dynamic optimisation based on trying a series of
'one-period optima based on positive portfolios only
Dim alloc_range As Variant 'not really a range
  'but do this to make consistent with static code
Dim partitions As Integer, sim As Integer
Dim noscensahead As Integer, scen As Integer
Dim i As Integer, nostrats As Integer, dummy As Double
Dim notargets As Integer, minsumsqdev As Double
Dim sumsqdev As Double
Dim S As Return_Index, oldS As Return_Index, j As Integer
Dim realreserve As Double, realsurplus As Double
Dim osterm As Integer
partitions = 3
dummy = Ran0(seed)
alloc_range = pentatope(partitions)
nostrats = UBound(alloc_range)
ReDim alloc_array(1 To nostrats) As Asset_Mix
notargets = target_range.Rows.Count
ReDim target(1 To notargets) As Double
ReDim asset_share(1 To notargets) As Double
ReDim stratbest(1 To notargets) As Integer
ReDim outvec(1 To notargets, 1 To 2) As Double
ReDim sumreturn(1 To notargets) As Double
ReDim sumsquare(1 To notargets) As Double
For i = 1 To nostrats
   With alloc_array(i)
      .Equity = alloc_range(i, 1)
      .Fixed_Interest = alloc_range(i, 2)
      .Cash = alloc_range(i, 3)
      .Index_Linked = alloc_range(i, 4)
      .Property = alloc_range(i, 5)
```

```
   End With
Next i
For i = 1 To notargets
   target(i) = target_range(i)
Next i
Select Case Model
Case "Wilkie"
   noscensahead = 10
Case "Chaotic"
   noscensahead = 1
Case "RandomWalk", "Fractal"
   noscensahead = 8
Case "Dyson&Exley", "JumpEq"
   noscensahead = 6
End Select
ReDim testreturn(1 To noscensahead, 1 To nostrats) _
   As Double, testRPI(1 To noscensahead)
ReDim vrealdisc(1 To noscensahead) As Double
For i = 1 To notargets
   sumreturn(i) = 0
   sumsquare(i) = 0
Next i
For sim = 1 To nosims
   For i = 1 To notargets
      asset_share(i) = 0
   Next i
   Unit S
   Do
      oldS = S
      'calculate hypothetical prospective returns for
      'scenario scen and strategy j
      For scen = 1 To noscensahead
         S = oldS
         Update S, Model, scen
         testRPI(scen) = S.Retail_Price
         vrealdisc(scen) = 1 / (1 + S.Real_Yield)
         For j = 1 To nostrats
            With alloc_array(j)
               testreturn(scen, j) = (.Equity * S.Equity _
                  / oldS.Equity + .Fixed_Interest * _
                  S.Fixed_Interest / oldS.Fixed_Interest _
                  + .Cash * S.Cash / oldS.Cash _
                  + .Index_Linked * S.Index_Linked _
                  / oldS.Index_Linked _
```

```
            + .Property * S.Property / oldS.Property)
        End With
    Next j
Next scen
osterm = horizon - S.Year
'now determine best strategy for each target
For i = 1 To notargets
   minsumsqdev = 1E+100 'any very large number will do
   'go through strategies one by one
   For j = 1 To nostrats
      sumsqdev = 0
      For scen = 1 To noscensahead
         realreserve = target(i) * _
            vrealdisc(scen) ^ osterm - _
            (1 - vrealdisc(scen) ^ osterm) _
            / (1 - vrealdisc(scen))
         realsurplus = (asset_share(i) + _
            oldS.Retail_Price) * _
            testreturn(scen, j) / testRPI(scen) _
            - realreserve
         If Model = "Fractal" And realsurplus > 0 Then
            'no credit for outperformance
            realsurplus = 0
         End If
         If Model = "Chaotic" Then
            'this is a fudge. Want to end up
            'choosing highest return
            sumsqdev = sumsqdev - realsurplus
         Else
            sumsqdev = sumsqdev + realsurplus ^ 2
         End If
      Next scen
      'see if this is the best yet
      If sumsqdev < minsumsqdev Then
         minsumsqdev = sumsqdev
         stratbest(i) = j
      End If
   Next j
Next i
'now full steam ahead with chosen strategy
S = oldS
Update S, Model
For i = 1 To notargets
   asset_share(i) = asset_share(i) + _
```

```
            oldS.Retail_Price
         With alloc_array(stratbest(i))
            asset_share(i) = asset_share(i) * ( _
               .Equity * S.Equity / oldS.Equity _
               + .Fixed_Interest * _
               S.Fixed_Interest / oldS.Fixed_Interest _
               + .Cash * S.Cash / oldS.Cash _
               + .Index_Linked * S.Index_Linked _
               / oldS.Index_Linked _
               + .Property * S.Property / oldS.Property)
         End With
      Next i
   Loop Until S.Year = horizon
   For i = 1 To notargets
      Select Case Model
      Case "Fractal", "Chaotic"
         If asset_share(i) > target(i) * S.Retail_Price Then
            asset_share(i) = target(i) * S.Retail_Price
         End If
      Case Else
      End Select
      sumreturn(i) = sumreturn(i) + asset_share(i) _
         / S.Retail_Price
      sumsquare(i) = sumsquare(i) + (asset_share(i) _
         / S.Retail_Price) ^ 2
   Next i
Next sim
For i = 1 To notargets
   outvec(i, 1) = Sqr((sumsquare(i) - sumreturn(i) ^ 2 / _
      nosims) / (nosims - 1))
   outvec(i, 2) = sumreturn(i) / nosims
Next i
PosRiskReturn = outvec
End Function

Function GearRiskReturn(Model As String, _
   horizon As Integer, seed As Double, nosims As Integer _
   , target_range As Variant) As Variant
   'mean variance optimisation
   'allowing for geared solutions
Dim notargets As Integer, noscensahead As Integer
Dim vrealdisc As Double, dummy As Double
Dim PVprems As Double, PVtarget As Double
Dim scen As Integer
```

```
Dim sim As Integer, S As Return_Index
Dim oldS As Return_Index
Dim W(1 To 5, 1 To 5) As Double
Dim U(1 To 5, 1 To 3) As Double
Dim osterm As Integer
Dim oldU As Variant, totU(1 To 5) As Double
Dim i As Integer, j As Integer
Dim returnvec(1 To 5) As Double
notargets = target_range.Rows.Count
ReDim target(1 To notargets) As Double
ReDim asset_share(1 To notargets) As Double
ReDim alloc(1 To notargets) As Asset_Mix
ReDim outvec(1 To notargets, 1 To 2) As Double
ReDim sumreturn(1 To notargets) As Double
ReDim sumsquare(1 To notargets) As Double
dummy = Ran0(seed)
For i = 1 To notargets
   target(i) = target_range(i)
   sumreturn(i) = 0
   sumsquare(i) = 0
Next i
Select Case Model
Case "Wilkie"
   noscensahead = 10
Case "Chaotic"
   noscensahead = 1
Case "RandomWalk", "Fractal"
   noscensahead = 8
Case "Dyson&Exley", "JumpEq"
   noscensahead = 6
End Select
For sim = 1 To nosims
   For i = 1 To notargets
      asset_share(i) = 0
   Next i
   Unit S
   Do
      'credit asset-shares with premium
      For i = 1 To notargets
         asset_share(i) = asset_share(i) + S.Retail_Price
      Next i
      'determine mean and variance-covariance matrix
      oldS = S
      For i = 1 To 5
```

```
        U(i, 1) = 0
        U(i, 2) = 0
        For j = 1 To 5
            W(i, j) = 0
        Next j
    Next i
    U(1, 3) = S.Equity
    U(2, 3) = S.Fixed_Interest
    U(3, 3) = S.Cash
    U(4, 3) = S.Index_Linked
    U(5, 3) = S.Property
    For scen = 1 To noscensahead
        S = oldS
        Update S, Model, scen
        osterm = horizon - S.Year
        'calculate everything in real terms
        returnvec(1) = S.Equity / S.Retail_Price
        returnvec(2) = S.Fixed_Interest / S.Retail_Price
        returnvec(3) = S.Cash / S.Retail_Price
        returnvec(4) = S.Index_Linked / S.Retail_Price
        returnvec(5) = S.Property / S.Retail_Price
        vrealdisc = 1 / (1 + S.Real_Yield)
        PVprems = (1 - vrealdisc ^ osterm) / _
            (1 - vrealdisc)
        PVtarget = vrealdisc ^ osterm
        For i = 1 To 5
            U(i, 1) = U(i, 1) + PVtarget * returnvec(i)
            U(i, 2) = U(i, 2) + PVprems * returnvec(i)
            For j = 1 To 5
                W(i, j) = W(i, j) + returnvec(i) * _
                    returnvec(j)
            Next j
        Next i
    Next scen
    oldU = U
    gaussj W, 5, U, 3
    For i = 1 To 3
        totU(i) = 0
        For j = 1 To 5
            totU(i) = totU(i) + oldU(j, 3) * U(j, i)
        Next j
    Next i
    For i = 1 To 5
        U(i, 3) = U(i, 3) / totU(3)
```

```
         U(i, 1) = U(i, 1) - totU(1) * U(i, 3)
         U(i, 2) = U(i, 2) - totU(2) * U(i, 3)
      Next i
      'The boring bit is done
      'Now determine asset allocations
      For i = 1 To notargets
         With alloc(i)
            .Equity = target(i) * U(1, 1) - _
                U(1, 2) + asset_share(i) * U(1, 3)
            .Fixed_Interest = target(i) * U(2, 1) - _
                U(2, 2) + asset_share(i) * U(2, 3)
            .Cash = target(i) * U(3, 1) - _
                U(3, 2) + asset_share(i) * U(3, 3)
            .Index_Linked = target(i) * U(4, 1) - _
                U(4, 2) + asset_share(i) * U(4, 3)
            .Property = target(i) * U(5, 1) - _
                U(5, 2) + asset_share(i) * U(5, 3)
         End With
      Next i
      'now this is for real
      S = oldS
      Update S, Model
      For i = 1 To notargets
         With alloc(i)
            asset_share(i) = .Equity * S.Equity _
               + .Fixed_Interest * S.Fixed_Interest _
               + .Cash * S.Cash _
               + .Index_Linked * S.Index_Linked _
               + .Property * S.Property
         End With
      Next i
   Loop Until S.Year = horizon
   For i = 1 To notargets
      sumreturn(i) = sumreturn(i) + asset_share(i) _
         / S.Retail_Price
      sumsquare(i) = sumsquare(i) + (asset_share(i) _
         / S.Retail_Price) ^ 2
   Next i
Next sim
For i = 1 To notargets
   outvec(i, 1) = Sqr((sumsquare(i) - sumreturn(i) ^ 2 _
      / nosims) / (nosims - 1))
   outvec(i, 2) = sumreturn(i) / nosims
Next i
```

```
GearRiskReturn = outvec
End Function

Sub gaussj(A, n, B, m)
'Gauss-Jordan elimination. Replaces B by A^-1*B
'and A by A^-1. A is an NxN matrix, and B is NxM
Dim i As Integer, j As Integer, k As Integer
Dim L As Integer, LL As Integer
Dim big As Double, irow As Integer, icol As Integer
Dim dum As Double, pivinv As Double
ReDim ipiv(1 To n) As Integer, indxr(1 To n) As Integer
ReDim indxc(1 To n) As Integer
For j = 1 To n
   ipiv(j) = 0
Next j
For i = 1 To n
   big = 0
   For j = 1 To n
      If ipiv(j) <> 1 Then
         For k = 1 To n
            If ipiv(k) = 0 Then
               If Abs(A(j, k)) >= big Then
                  big = Abs(A(j, k))
                  irow = j
                  icol = k
               End If
            ElseIf ipiv(k) > 1 Then
               MsgBox "Singular Matrix"
               Exit Sub
            End If
         Next k
      End If
      Next j
      ipiv(icol) = ipiv(icol) + 1
      If irow <> icol Then
      For L = 1 To n
         dum = A(irow, L)
         A(irow, L) = A(icol, L)
         A(icol, L) = dum
      Next L
      For L = 1 To m
         dum = B(irow, L)
         B(irow, L) = B(icol, L)
         B(icol, L) = dum
```

```
      Next L
    End If
    indxr(i) = irow
    indxc(i) = icol
    If A(icol, icol) = 0 Then
      MsgBox "Singular Matrix"
      Exit Sub
    End If
    pivinv = 1 / A(icol, icol)
    A(icol, icol) = 1
    For L = 1 To n
      A(icol, L) = A(icol, L) * pivinv
    Next L
    For L = 1 To m
      B(icol, L) = B(icol, L) * pivinv
    Next L
    For LL = 1 To n
      If LL <> icol Then
        dum = A(LL, icol)
        A(LL, icol) = 0
        For L = 1 To n
          A(LL, L) = A(LL, L) - A(icol, L) * dum
        Next L
        For L = 1 To m
          B(LL, L) = B(LL, L) - B(icol, L) * dum
        Next L
      End If
    Next LL
Next i
For L = n To 1 Step -1
    If indxr(L) <> indxc(L) Then
      For k = 1 To n
        dum = A(k, indxr(L))
        A(k, indxr(L)) = A(k, indxc(L))
        A(k, indxc(L)) = dum
      Next k
    End If
Next L
End Sub
```

APPENDIX D

GLOSSARY

Appraised value	The result of a discounted cash flow calculation, which is intended to be consistent with a market basis.
Adapted	A strategy which can be carried out based only on information available at the decision date. Sometimes called *non-anticipating*.
Added value	Increase in economic value.
Arbitrage	A trading activity which aims to generate high returns with zero risk. This is often believed to be impossible, leading to the construction of *arbitrage free* models.
Arbitrage pricing theory	A model for asset prices based on the absence of arbitrage. This is essentially a multi-dimensional version of the CAPM.
Ask value	The lowest price at which a market maker has advertised his willingness to sell a security (or, equivalently, the price at which anyone else can buy a small quantity).
Assessed value	The value of an asset as used in a pension fund valuation on an ongoing basis. Calculated by discounted cash flow, and not, as a rule, equal to market value.
Asset-liability study	An investigation into the assets and liabilities of a financial entity, including the interactions between them.
Autoregressive conditional heteroscedasticity (ARCH)	A data series exhibits the ARCH effect if periods of high volatility alternate with periods of lower volatility.
Autoregressive model	A model in which the direction of future movements may depend on past movements and the current position.
Bid value	The highest price at which a market maker has advertised his willingness to buy a security (or, equivalently, the price at which anyone else can sell a small quantity).
Black-Scholes formula	A formula for pricing options based on an interest rate, the value, yield and volatility of the underlying asset.
Brownian motion	see *random walk*

Budget constraint	The constraint imposed by the initial amount of cash available for investment.
Capital asset pricing model (CAPM)	A particular model which describes a relationship between the risk of an investment and the expected return.
Chaos	A mathematical model in which future movements are highly sensitive to initial conditions.
Complete market	A market in which a market value can be observed for every cash flow stream.
Concave function	A function such that the region below it is convex; a function which looks like $\cap$, (or).
Convex function	A function such that the region above it is convex; a function which looks like $\cup$, (or).
Cost of capital adjustment (COCA)	An adjustment to the present value of a series of cash flows to allow for the cost of servicing the risk capital backing the business.
Deflator	A stochastic generalisation of discount factors, of importance in dynamic optimisation.
Derivative	A security whose cash flows derive from the price of another underlying asset.
Diffusion model	A model in which random variables are a continuous function of time. An example is Brownian motion.
Dynamic hedging	Continued readjustment of investment exposure, with the objective of reducing risk from derivative positions.
Dynamic optimisation	Optimisation of a strategy assuming that any decision may be reviewed at annual or more frequent intervals.
Earned rate	The expected (or best estimate) return on assets held.
Economic value	Can mean various things; usually either market value or marginal value.
Efficient frontier	The set of efficient portfolios.
Efficient market	A market in which there are no trading opportunities.
Efficient portfolio	An investment portfolio which, for a given expected return, has the lowest risk, or for a given risk has the highest expected returns.
Empirical law	The probability law which is believed to describe the real world, based either on historical data or on encoding subjective views of the future. Sometimes called the *true law*.

Equivalent martingale measure	The same as a *risk neutral law.*
Fenchel conjugate	An intermediate stage in dynamic optimisation, which enables the impact of the market to be analysed separately from the utility of the particular investor.
Financial economics	The application of economic theory to financial markets.
Fractal	A geometric construction where large parts of the structure can be mapped onto smaller parts of the same structure by a change of scale.
Ho-Lee model	One model describing the term structure of interest rates and associated derivatives.
Jump process	A model in which economic quantities are not continuous functions of time.
Local time	The amount of time spent by a time dependent random variable at a given level.
Locking in adjustment	A downward adjustment to the value of free assets to reflect the fact that these assets are subject to additional risk from the underlying business. Economically equivalent to a *cost of capital adjustment.*
Long-term value	Synonym for *assessed value.*
Long-term return	The limit of the average return measured over a long period. Actuaries usually mean a geometric average.
Marginal value	A cash flow today which would fairly compensate an investor for a small quantity of an asset.
Market value	The price at which an asset trades in the market.
Mean-variance analysis	Identification of the efficient frontier, using variance as a measure of risk.
Monte Carlo simulation	see *simulation.*
PEP	Abbreviation for *personal equity plan*, a form of tax efficient equity investment for private individuals.
Preference relation	A description of a rule which determines whether a market participant will prefer one cash flow stream to another.
Present value (PV)	The result of a discounted cash flow calculation.
Quadratic programming	The optimisation of a quadratic function subject to linear constraints.
Random walk	A model in which future movements from the current value are independent of past movements. It is sometimes also required that the future movements should have a normal or Gaussian

	distribution, in which case we have a Gaussian random walk or, in continuous time, a Wiener process or Brownian motion.
Rational expectation hypothesis	The belief that the current yield curve contains implicit forecasts of the path of future interest rates.
Risk discount rate	An interest rate for use in discounted cash flow calculations, taking into account the riskiness of the cash flows to be discounted.
Risk-free rate	A theoretical discount rate useful in discounted cash flow analysis, particularly in the context of a risk-neutral law.
Risk-neutral law	An artificial probability law under which the expected return on all assets is equal, irrespective of risk. This is distinct from the *empirical law* which describes the true probability distributions.
Risk return plot	A chart showing various portfolios with risk on the horizontal axis and expected return on the vertical axis.
Sample path	A single economic scenario (possibly describing several economic quantities) from a Monte Carlo simulation exercise.
Shadow probability space	The same as a *risk neutral law*.
Shareholder value	This means different things to different people. I have used it as a synonym for *marginal value*.
Simulation (or Monte Carlo simulation)	The random generation of economic scenarios or sample paths from a stochastic model.
Static optimisation	Optimisation of a strategy on the basis that a single strategy will be implemented for all future time.
Stochastic model	A model describing the probabilities of various events.
Term structure	The relationship between the yield on a bond and its maturity.
Trading opportunity	A trading strategy which produces an abnormally high expected return given the level of risk.
Transaction costs	The cost of buying and selling investments. I have used this to mean the difference between bid and ask prices; some economists also include the effect of the trade moving the market as part of the transaction cost.
Value	This term on its own is used to mean many different things, and I have tried to avoid using it.

Volatility	A measure of how much a price is likely to change in a unit time. In option pricing, this usually means the standard deviation of the log price move over one year.
Wiener process	see *random walk*
Wilkie model	A stochastic model for actuarial use, describing various asset classes.

ABSTRACT OF THE DISCUSSION

Mr A. D. Smith (introducing the paper): I put to you five actuarial puzzles:

(1) This regards the choice of rates of interest for discounting liabilities when the assets are predominantly invested in equities. In life assurance, for the purposes of the DTI Returns, we discount liabilities at the running yield, while in pensions work we discount at an expected total return. In non-life insurance we often use gilt yields less a margin, irrespective of the assets held. What differences in the nature of the businesses lead to these different approaches being applied?

(2) This relates to the valuation of assets. According to the *Financial Times* of 20 March, pensions actuaries "often ignore market price as volatile and untrustworthy and construct other, more stable measures of long term value". The same idea has been proposed for insurance work, but was rejected by the profession in favour of a more market-based approach. In practice, for management purposes, some of the assets get sucked into a profit test, from which cash flows emerge and are discounted at an arbitrary rate; the asset valuation which results is closer to a discounted cash flow approach than to market value. If I were an actuary working overseas, I might value assets at book value or apply some arbitrary smoothing process to market returns. Why do we adopt so many different approaches to solve the same problem?

(3) This relates to the provision of guarantees. What is the cost to the estate of providing guarantees for with-profits funds? Why is it so capital intensive for insurers to offer unit-linked policies with maturity guarantees, compared to the same product offered by a bank?

(4) This relates to asset-liability modelling. Insurance work in this area is still largely based on immunisation theory, which requires gilt portfolios to be switched from time to time. Why do currently popular pension fund methodologies not allow funds to switch assets in an optimal way to respond to changes in investment conditions? Allowing for such switching behaviour would often lead to dramatically different advice being given.

(5) This relates to investment outperformance. Consultants earn healthy fees for fund manager selection; a process which should lead to out-performance relative to the prescribed benchmark. Why then do we not allow for this out-performance when constructing a model office or calculating appraised values?

At its most basic, actuarial science consists of a tool kit of algorithms. These algorithms have been adapted to different practice areas. In most cases current practice is highly fragmented, reflecting the adaptations favoured by particularly eminent individuals working in each specialised area. This gives rise to inconsistencies, some of which I have tried to identify in my puzzles. In defence of the status quo, we may observe that no major financial catastrophe has yet been conclusively pinned on the failure of actuarial formulae. However, if I scratch the surface of the actuarial method, I am hard pressed to find a unifying theory.

Many of our counterparts in the accountancy, banking and management consultancy professions have embraced financial economics more enthusiastically than the actuarial profession has. On the other hand, active fund managers are notoriously sceptical. It would be quite legitimate to hold back from financial economics if we had serious doubts about the validity of the framework it provides. Most of us accept that we need to make assumptions. If we do not like others' assumptions, there is nothing to stop us from re-applying the techniques based on our own assumptions. However, the greatest restraint on the profession's adoption of these new techniques is the impenetrable forest of advanced mathematics and jargon which greets any lay reader of the financial economics literature. It is not immediately evident how this mathematics is to be applied; the formulae are not generally of the kind which enable simple assumptions to be put in or straightforward deductions to be drawn. Ironically, among actuaries who, according to popular perception, specialise in impenetrable mathematics and obscure jargon, few have shown the tenacity necessary to turn the theory into workable practical tools.

I have built models using financial economics and worked through applications in grinding detail, from the initial assumptions to numerical output, going down to the level of line-by-line listing of the computer programs I have employed. The process of producing this work was highly educational for me. This has resulted in a highly technical paper, and I believe that this is what is needed in order for the profession to get to grips with this new and exciting area.

Has financial economics helped me to solve my five puzzles? Frankly, they still remain puzzling. Conventional financial economics suggests that market values are key, and that discounting ought to relate to the expected total return on matching assets. The banking approach to options and guarantees is probably preferred to that recommended by the actuarial Maturity Guarantees Working Party. Asset-liability modelling probably should allow for asset switches. There does not seem to be any satisfactory economic theory of investment skill; some economists would even suggest that fund manager selection is a waste of time in principle.

Some of these answers may not seem satisfactory, but the underlying economic framework is internally consistent, simultaneously describing current market prices, future movements and investor behaviour; this is the benchmark against which other professionals will judge us. It would not be difficult for any numerate professional familiar with financial economics to come to similar conclusions regarding the inconsistency of current actuarial practice.

I commend financial economics to the actuarial profession today. The consequences of adoption would be to open new vistas of productive work and an enhanced ability to communicate with other financial professionals, at a time when we are anxious to expand the profession's horizons into wider fields.

Mr A. J. Wise, F.I.A. (opening the discussion): On 22 March 1993, at the Sessional Meeting, we held a debate where the motion proposed was that 'This house believes that the contribution of actuaries to investment could be enhanced by the work of financial economists'. The author of this paper spoke in support of the motion, which was, in the event, carried. Clearly inspired by his theme, the author, in this paper, has explained to us how he thinks actuaries can use financial economics, not just in investment, but in all areas of our work. I am sure that the unfamiliar mathematics will have deterred many actuaries, but this is a paper which has both breadth and depth.

Section 2 is entitled 'The Concept of Value', and the first paragraph identifies several distinct concepts of value. There is mention of present value, economic value, appraised value, assessed value and actuarial reserves. The author then provides a commendably simple reconciliation of three different value concepts, using for his example a cohort of single premium PEP business, in which a value is to be placed on the expected future profit stream. The author concludes, in ¶2.5.4, that valuation, either by risk discount rate, or by adjusting for the cost of capital, or by discounting risk neutral expectations, is equivalent, and that any debate about the merits of the methods is one of presentation rather than economic substance. This, it seems to me, is an important statement which seeks to align and consolidate the separate thought processes within actuarial work and financial economics.

Reverting to the author's distinction between different concepts of value, I saw that he did not mention the term 'actuarial value' or try to give it a meaning. I then asked myself: does the actuarial profession not put its own name to a clearly defined concept of value? If not, that is surely a serious omission. As we move into the 21st century our profession will need to maintain its credibility with business managers and senior accountants, most of whom will be graduates well versed in the vernacular of financial economics. If we do not share their vernacular, then they will not relate to our philosophy and techniques.

At present it is still commonplace for actuarial reports on pension funds to quote an 'actuarial value' of the assets, which differs from market value. How much longer can we go on quoting actuarial values as if they have some meaning in the real world, when these figures can be inconsistent with economic value as measured by reasonably efficient markets on a particular date? If our profession is to make progress, should we not, at least, agree a common definition of actuarial value for all purposes? Should not that definition be: "the result of a discounted cash flow calculation which is consistent with economic value"? We could then use this expression in place of the rather vague sounding 'appraised value'. We could also bring the presentation of pension fund valuations of assets and liabilities into line with market conditions at a given time. The key to this transformation is to acknowledge that the value of the liabilities is the value of the appropriate assets which are required to meet those liabilities.

In ¶2.3.7 the author refers to the current debate regarding transfer values into and out of final salary pension schemes. A more far-reaching point is the choice of discount rate in relation to accounting for pension costs in company accounts. The United States standard FAS87 effectively requires discounting at the risk-free rate associated with bonds. The corresponding United Kingdom standard SSAP24 is now being revised by the Accounting Standards Board, whose 1995 discussion paper argues in favour of discounting at the expected rate of return on assets. With the high equity backing of most U.K. final salary pension schemes, the U.K. approach implies the use of a higher risk-adjusted discount rate, which brings it into potential conflict with the principles of FAS87. Now the International Accounting Standards Committee has begun to look at this issue, and, because of the globalisation of capital markets, there is a clearly perceived need for rapid global standardisation of accounting principles.

The principle for deciding the pension liability discount rate must be established one way or the other, and a working party of our Pensions Board is currently looking at the problem and discussing it with the accountants. For my part, I am clear that a company should be assessing its final salary pension cost using a risk-adjusted discount rate, having regard to an investment policy which is appropriate to the nature and risks of the liabilities. It seems, however, that the accountants are less clear about this, and that we may have some difficulties in persuading them that the U.S. standard is flawed. This paper highlights for me the potential strength which can be exerted in favour of the U.K. approach if we appeal to the fundamental principles of financial economics. Thus, when the accountants suggest that a risk-free discount rate should be used, we should respond by saying that the non-diversifiable risk of pension liabilities indicates a higher risk-adjusted discount rate. When the accountants suggest that a lower discount rate is justified on grounds of prudence, we can, I think, point out that any risky preference debt of a company is effectively valued in the balance sheet under FRS4 at a net proceeds value which is equivalent to discounting at the risk-adjusted rate. Financial theory then tells shareholders how to value their shares to allow for the risky liabilities as valued this way without a margin of prudence in the company's balance sheet.

In Section 2.5 the author introduces the concept of risk-neutral valuation. He explains that any sensible market pricing of assets and liabilities is necessarily consistent with an actuarial valuation of expected future cash flows, where the expectation of uncertain future cash flows is derived from an appropriate and self-consistent basis of probabilities which is, however, shifted away from best estimates. The shifted probabilities are those which an investor would need to believe in order to justify actual market prices, but without requiring any risk premium. The probabilities are, therefore, those of a risk-neutral investor. Any market value is an actuarial value as seen by a risk-neutral investor. This is surely a profound idea. At the very least it is reassuring to be told that an actuarial valuation is important and fundamental in any coherent system of pricing assets and liabilities of whatever nature.

I first encountered the risk-neutrality result a few years ago, and, in order to understand why it is true, I referred to Harrison & Kreps (1979). However, I found it impossible to understand! I then worked through the thinking shown in Section 2.6 concerning option pricing. Whilst I agree with the author's comments in this section, I see no point in complicating the presentation, as he does, by using the dividend adjusted Black-Scholes formula. Instead of complicating things with dividend adjustments, why not make the exposition simpler?

The theory of option pricing is woefully lacking from the actuarial education syllabus. Many of the text books use a simple binary model to explain option pricing, and I found that it helps to work through the binary model and relate it to actuarial valuation. The actuarial principle of matching is precisely equivalent to the hedging argument which was used by the financial economists to solve the option pricing problem.

Section 3 is a review of asset/liability methodology. I think that the author is agreeing that, whilst financial economists tend to work with short time horizons, actuaries generally work to long time horizons and need appropriate stochastic models to tackle some of the problems which arise.

In Section 3.2 four aspects of the economy are identified, which the author says should be covered by a complete model. Given the importance of such models in our present and future professional work, can we agree on these criteria? Are they necessary and sufficient? Should we, indeed, adopt an agreed set of principles such as these in order that the profession can exercise some degree of supervision over the creation and use of models?

The author goes on to investigate, in some detail, the characteristics of various published models, including those of Wilkie, and of Dyson & Exley, together with a model of his own design which is described as a jump equilibrium model. In working through his analysis of the various models, it seems to me that the author has pointed out a way forward. His analysis of the alternatives could be a paradigm for the way that models should be compared in future.

In ¶5.5.2 it says that different models, or at least different versions of a model, may validly be believed by different market participants, all of whom are looking at the same market prices. Is it then futile to search for a holy grail of the one true economic model? Should we, instead, be looking to identify a catholic variety of models which are both useful for actuarial purposes and consistent with economic reality?

The paper ranges over many areas of actuarial work: pensions, insurance and investment. The author has the confidence to apply his thinking to corporate finance generally. I think that he has presented a more coherent view of different areas of actuarial work than would be possible without the use of financial economics. He has certainly inspired me to think that, if we do not align ourselves with the language and the philosophy of financial economics, we risk losing our credibility and influence to convince others, such as the accounting standards bodies, of the strengths in our own philosophy.

Mr P. J. Lee, F.I.A: In ¶2.3.7, with regard to pension fund values, the cost of capital adjustment (COCA) should probably be less than the difference between equity and gilt-based transfer values. This is because the scheme member shares some of the risk with the employer, for two reasons: first, the fund sponsor's guarantee is not usually as credit-worthy as the government's; and secondly, there may well be some element of discretionary increases. This would justify using an equity return less, say 1% p.a., to pay for the COCA, rather than a gilt return which might be 2% p.a. below the equity return.

In ¶2.5.3 I like the concept that the market may be inefficient whilst still featuring absence of arbitrage. My own view is that markets are not yet efficient because of factors other than the obvious ones of tax and transaction costs:

(1) The majority of investors have short time horizons (typically much less than 5 years) and are very likely to bid up the price of bonds relative to equities. Historically, it has been rare for bond investors to outperform equity investors over longer time periods, even over calamitous periods including wars and depressions, probably because of the weight of short-term investors. It follows that investors with longer time horizons can outperform the market, providing that they remain in a minority.

(2) Peer pressure and the fear of being blamed ensure a strong herd instinct. This means that many investors will be reluctant to exploit perceived inefficiencies to a significant degree until everyone else starts doing so as well. This is very clearly exemplified by the reluctance of investors all round the world to diversify internationally, which, in many cases, is compounded by local legislative constraints.

With regard to ¶3.3.6 and the complaint that stationary stochastic models offer little insight relative to deterministic projections, the 'sufficiently long time horizon' that one has to project over is well over two hundred years before the risk parameters cease to have a significant influence over the results of stationary models.

With regard to ¶¶3.4.1 and 3.4.2 and the value of asset liability studies, I think that the author will agree with me that at least half of this value comes from the information provided with regard to the quantification of risk. Stochastic modelling remains the only way to obtain such quantification.

The cointegrated models are supposed to allow for instability in the economy, whilst retaining stability in relative returns. However, because of a technical point to do with the possibility of negative yields on all assets, the mean real returns on all assets under the author's model have to drift upwards over time, leading to real returns of the order of 15% p.a. over 100 years. For more typical periods of 20-50 years, this effect is less obvious, but it is still there. The upward drift of mean returns may make it difficult for the user to understand the impact of different input parameters into the model, and therefore to choose those parameters.

The author states, in ¶7.2, "All models are wrong in some respect, but in the real world we have to make decisions. Either we use a flawed model to help, or we use no model and rely on 'feel'". I would add that those who rely on feel are really choosing to use a model which is their own private one, unpublished and totally subjective.

Mr C. J. Exley, F.I.A.: I welcome the author's process of first establishing a pricing law and then fitting a model, rather than relying mainly on subjective ideas of 'realism' to fit a statistical model to past data and then trying to infer things from that. Pricing laws perform the role of conservation laws in physical science, since they put constraints on how things can develop over time. So the author's comparison, in ¶2.6.7, between building a financial model which violates the law of one price and building a perpetual motion machine is subtle and profound. Conservation laws distinguish subjective art from objective science and, possibly, also distinguish much actuarial folklore from financial economics.

Of course this more rigorous approach also places some constraints on what we can sensibly say as actuaries. For example, the law of one price says that: "if I take a short position of £100 in a corporate bond, the price of that short position (similar to a pension liability) is unaffected by where I invest my £100 proceeds (similar to pension fund assets)". In this pricing framework it is nonsensical to talk about an equity-linked price of this short position (if I buy £100 of equities) or a gilt-linked price (if I buy £100 of gilts). This sits rather uncomfortably with the idea of an equity-linked cash equivalent, which seems, at first sight, to confuse assets with liabilities quite spectacularly.

The usual escape from these constraints is to distinguish between value and price. The author has rebutted the idea that market prices are volatile relative to subjective actuarial values. However, he does not address the issue of consistency in a valuation. Here again the usual argument is easily reversed. A market basis is always consistent. On the other hand, at some point we always need to mix actuarial values with prices if we use an off-market basis. For example, we usually accept the current level of the Retail Prices Index as a starting point, despite the fact that these prices are also set by marginal agents and are therefore, presumably, as unreliable as asset prices as a base from which to project using long-term assumptions.

In connection with asset-liability modelling, the author demonstrates, in Appendix A, that optimisation over the long term can be collapsed into a series of short-term optimisation problems. Also, the author's dynamic solution shows that, for each short-term problem, we need to know what we might do tomorrow before we can decide on our optimal policy for today. This reintroduces long-term actuarial time horizons into the mathematics, but in a much more subtle and sophisticated way than if we take one giant leap to an ultimate horizon.

It is a reassuring validation to see that, forty years on, this approach suggests that an immunised strategy is the best match. It seems to me that giant leap models, which rely on supposed time diversification of risk, which argue, for example, that equities can 'match' any long-term liabilities you like, are unambiguously flawed.

Mr G. S. Finkelstein: A large part of financial economics is based on mean variance optimisation. This is discussed at some length in Section 3. In ¶3.3.4 the author mentions stable distributions as an alternative to the Normal distribution, which is used in large parts of financial economics. He points out that, apart from the Normal distribution, which is a special case of the more general stable family, all other stable distributions have infinite variance. In fact the log stable distribution also has infinite mean. Therefore, if one generalises the assumptions from the normal to the more general stable, and there are strong arguments for doing so, large parts of financial economics based on optimising means and variances are no longer applicable.

You could modify the theory by optimising, for example, medians and inter-fractile ranges in the place of the mean and the variance. However, you do not necessarily arrive at the same conclusions as those derived from optimising means and variances assuming Normal distributions. This is because stable non-Normal distributions can introduce more extreme stochastic fluctuations, but not as many less extreme fluctuations.

Mr J. Bridgwater (in a written contribution that was read to the meeting): My thoughts and comments reflect the importance of your work from a banking perspective.

A subset of the techniques described in the paper is already in consistent use within the banking world — particularly in relation to the derivatives business lines, but there is enormous scope for the banking world to learn from the application of financial economics to other, less familiar, problems. Quantitative specialists within the actuarial and banking worlds have independently developed sophisticated techniques for the analysis of pricing and risk questions, which have hitherto not overlapped to any great extent. Such a distinction is, in many respects, artificial, as is clearly demonstrated by this paper. The development of financial economics research can benefit both the banking and insurance worlds and promote synergies throughout, by allowing each to learn from the successes and strengths of the other.

The concept of pricing in the framework of a risk-neutral law is fundamental to most of the quantitative work carried out within banks. Any bank that is involved in trading options or more complex derivatives on underlying stocks, bonds, currencies, other commodities or interest rates will have developed a considerable armoury of numerical techniques for pricing complex structures in a Black-Scholes world. The use of the Black-Scholes assumptions allows a risk-neutral framework equivalent to a simple partial differential equation which is amenable to solution using a number of well-known techniques. Closed form solutions exist for many simple products.

The deficiencies of the Black-Scholes paradigm are well understood by practitioners, but the model itself still provides considerable value, by reducing the pricing problem to the identification of a simple hedge portfolio comprising traded assets. The remaining model error can be compensated for by using an 'implied volatility' — implied from the option's market price. For interest rate derivatives, in particular, more advanced models have been developed and implemented, using extensions to the numerical techniques developed for more simple problems. Other extensions to the Black-Scholes world try to allow for observed departures from the model assumptions, such as non-constant volatility or mean reversion of interest rates. The development of more advanced models by banks has been hindered, in many cases, by a general lack of understanding of the underlying financial economics, leading to a great reliance upon academic work. In particular, many of the currently fashionable interest rate models are derived from academic papers.

The extension from Black-Scholes to more sophisticated models requires careful consideration of a number of issues that are not immediately intuitive or tractable. The paper contains a lucid and detailed exploration of some of these issues, and a number of the results that the author has obtained will be of significance beyond the world of insurance.

The development of a jump equilibrium model, in particular, is of significant value in my speciality field — foreign exchange (FX) derivatives — where a combination of diffusion and jump processes would appear to be a very close model of the observed reality. Indeed, a simple jump-diffusion model, mentioned in the paper (the cost of capital adjustment model), has been successfully calibrated to observed FX option prices, FX being the most liquid actively traded derivatives market.

Many of the derivatives markets actively traded by banks have reached a degree of maturity and liquidity that renders questions of market completeness relatively unimportant. For that reason, amongst others, this area of research has been largely neglected by the banking world, even though the insights available are stimulating, and for emerging market places are important. As the derivatives businesses expand into new markets, in search of greater margins and wider client coverage, questions of illiquidity, transaction costs and the impact of external controls will assume a larger significance to the quantitative analyst. These are areas where actuarial expertise has significant value to add.

Increasing liquidity in trading markets also leads to an increased interest as to the capital requirements for the trading business. This analysis is currently being performed by the banking world under titles such as 'Value at Risk' or 'Risk Metrics'. It is very clear from this paper that the questions to be answered are equivalent to those posed by actuaries for life assurance businesses.

The use of a conditional expectations approach to pricing claims on untraded (or untradeable) risks coincides neatly with an approach recently developed in our bank, whereby many complex multi-factor claims are priced by calculating a conditional expectation of the payoff so as to reduce the dimensionality of the problem.

There is clearly considerable benefit to be gained for the quantitative profession from closer links between its actuarial and banking arms. I hope that the publication of research of this quality can lead to the further development of such links in the future.

Mr J. Plymen, F.I.A.: In ¶6.2.3 the author refers to remarks in Clarkson & Plymen (1988). In our paper we said "we are satisfied that the Markowitz-Sharpe model using historic betas for return and risk has no practical application to portfolio management, either to improve performance or to reduce the risk". The remarks made in 1988 were based on the strong evidence that we were able to produce at the time, and were explained in our paper. We stand by this statement.

Since 1988 I have discovered a most important item of further evidence which supports our cause. It is all a question of the way we calculate the risk. The principle is to maximise the expected yield, allowing for the risk. The expected yield is derived from investment analysis. It allows for all the possible adverse and favourable factors of the company concerned, and so the risk has been allowed for. On the other hand, Markowitz, in his system, obtains the expected yield or the return by study of price and dividend history. He works out the beta factors, taking the movement of prices and dividends monthly over a period.

I ask the author a direct question: does the Markowitz system of deriving a return based on historic factors produce an answer before the risk or after the risk? This question is most important. If the Markowitz system produces a before-risk result, then it is all right to deduct the risk and get the net return.

On the other hand, if the Markowitz system is producing an after-risk result, then we have to conclude that the risk factor has already been allowed for in calculating the expected return, so there is no scope for making any further deduction for the risk. In these circumstances the whole mean variance procedures are open to question. The mean variance procedures involve an analysis of the risk item. If the risk item being deducted from the return is nil, then you cannot analyse it.

Mr M. Cooper (a visitor): I am a derivatives trader and a chartered accountant. As a derivatives trader I have considerable problems interpreting the mathematics of this field. Consequently I will restrict my comments to a more qualitative assessment of some of the techniques involved at a very high level.

It is certainly true that the way that derivatives traders use financial economics is by using the models to make a transformation between a non-traded derivative (either an exotic option or some other structure which is not traded in the market) and the related traded securities, both for the purposes of pricing and hedging. In the same way, for example, that a yield to maturity can be used to compare two different bonds of notionally different coupons and prices, so we use an option pricing model, such as Black-Scholes, to compare options of different characteristics. Thus, I can use the implied volatility of a traded option to price an option of widely differing characteristics, and, far more importantly from my perspective, I am also able to generate a hedge cluster. Thus, I am concerned with the relative price of those two items; in a sense I am using a controlled variate technique to price one from the other, and provided that they both price at the same level, then I am quite happy.

Other applications of financial economics in banking perhaps have more application to areas of actuarial interest — in particular the uses being made of value at risk and risk advisory projects. Considerable success has been achieved in applying the same models that we have developed for trading purposes to the more fundamental questions of evaluating a client's portfolio, a client's position and making estimates of the maximum potential risk to which the portfolio is subject.

The jump equilibrium model that is proposed by the author is of great interest to us, and is a method of tackling some of the major simplifying assumptions that are common to most of the other models that are currently being used.

Professor R. S. Clarkson, F.F.A.: The author and I agree that the current theories of financial economics are seriously incomplete. We disagree on whether they are dangerously unsound. The author said that if you do not like some of the assumptions then you can change them, but a number of speakers, including Mr Cooper, who is at the 'sharp end', referred to what the author called the 'impenetrable forest of mathematical jargon'. Much of the research that I have carried out over the past year is trying to find my way through this impenetrable jargon, and I have concluded that, in many ways, the financial economists are using the wrong type of mathematics. Others have come to similar conclusions.

If you look at the independence axiom which is at the key of linearity or non-linearity, the most vociferous opponent was Milton Friedman. Also, the American axiomatic approach to utility theory was criticised by Maurice Allais, who won the Nobel Prize for economics in 1988, two years before Markowitz, Miller and Sharpe were jointly awarded the Noble prize for their work in financial economics.

A number of the observations that the author makes on the basis of his practical experience support my own conclusions. In particular, he concludes that we need to make explicit allowance for the potential irrationality of others, and that simplistic linear models can be a recipe for what I call highly unintelligent behaviour.

The author has suggested that the models built by financial economists are now much less unrealistic than previously, in that many of the restrictive assumptions have been relaxed in recent years. For the moment we must differ on that, but we are not in disagreement as to the general way forward. We take what is there; we look to see if we are unhappy with some of the simplifying assumptions and we then try to build something better.

Mr K. P. W. Larner, F.I.A. (in a written contribution that was read to the meeting): My comments arise from a general insurance perspective; I will be particularly interested if actuaries in other specialist areas disagree.

The sentiments of the early sections of the paper are to equate appraised values with economic values. Appraised values are described as calculations of the present value of future cash flows, at appropriate discount rates, and economic values are equated with a market trade price. In the context of buying and selling companies, or indeed swapping cash flow streams by some other means of transfer, these sentiments are clearly right.

I would go further and say that, when an appraised value calculation using appropriate risk discount rates differs from another party's view of economic value of the same potential cash flow streams, then opportunities arise for arbitrage. On this basis, almost all actuarial projections and present value calculations have to be thought of in the economic conceptual framework. Thus I disagree with the paper, in ¶1.5, when it says that we should wait until "the profession reaches consensus on method" before starting to think "about how it is to be presented to the outside world". We cannot wait. Every day every actuary is typically involved in this process.

I believe that getting the conceptual framework right is probably more important than carrying out the detailed mathematical calculations. We need to work on our modelling techniques, and I believe that the paper is extremely useful in the examples it gives, but bigger errors can arise from wrong thinking. In the context of general insurance, I commend two papers that give a full description of the ways in which general insurance actuaries think: Bride & Lomax (1994) and 'The Valuation of General Insurance Companies' by J. P. Ryan & K. P. W. Larner (*J.I.A.* **117**, 597–669). I was struck by the commonality of the approach, conceptual thinking and language of these two papers and the one being discussed now.

The definitions of appraised values in ¶¶2.1.1 and 2.2.1 are not quite consistent. I see no reason why an appraised value, from the point of view of a proprietor, used as a management analysis tool, should necessarily take the viewpoint of a market trader.

Paragraph 2.5.2 is very important from a conceptual point of view, and I would welcome some expansion of it by the author. There are a number of different concepts here which make matters simpler if they are well understood:

(1) The first point is about unbundling the problem. When an actuary performs a present value calculation on some cash flow, he or she generally thinks of a range of possible outcomes for the cash flow. These could be modelled by simulations, and this is what much of the rest of the paper is based on. An example of this thought process is when one company buys another. The buyer is, in fact, buying all the possible outcomes of these cash flows. This includes the potential value at which the company which has been bought might be sold at a future date. In a simulation study, each of these cash flows is given an equal probability, some cash flows are generated by holding onto the organisation purchased and some generated by selling it at a future date. The conceptual mathematical framework on which a valuation of this unbundled set of cash flows is made can vary, hence the 'shadow probability space' concept, the 'equivalent martingale measure' concept and the 'risk-neutral probability law' concept. The important point appears to be the ingredients. These are the unbundled set of potential future cash flows, the concept of time value of money and equating to the price paid for, or the value put on, the unbundled set of cash flows.

(2) My understanding is that, essentially, linearity arises because you can sell parts of the organisation or its cash flows, or at least cost the out-sourcing of various services. Therefore, the ability to swap cash flows in and out of the unbundled set on a risk-neutral basis implies linearity.

(3) The use of the expected value of cash flows is purely for mathematical convenience, and, in a sense, part of the definition of the risk discount rate.

Taking this approach, I am interested in others' views as to whether actuaries and those in specialist areas regard their estimates as mean values, or are they just the most likely outcome or mode values? Generally I prefer to use the term 'central estimate of value', as it does not carry connotations of a probability distribution, except that it is somewhere near the middle.

Uncertainty may arise from the quality of data in investigations or from genuine randomness of the various processes. Separating these issues is a crucial part of actuarial work, and is worth explaining. If we really should be using expected values, then we will probably need to rethink our methodologies.

In much of the paper capital is described as investments put at risk. I prefer to view capital more as a guarantee. This is exactly the point, in ¶3.1.4.1, that "capital structure is irrelevant", provided that the security risk is considered. There are many general insurance vehicles that operate in part without explicit cash or other investments as capital, but on the basis of guarantees. These would include Lloyd's, mutuals, P & I Clubs and, to a large extent, captives. Part of the capital is the ability to make future calls.

The example in ¶2.3.7 regarding transfer values I found interesting in this context. The issue of cost of capital adjustment really relies on who bears the risk. Since the sponsoring company may cease to contribute or go insolvent, at least part of the risk must fall on the members of the scheme.

Professor D. S. Wilford (a visitor): One of the basic tenets that runs throughout this discussion is our basic acceptance or non-acceptance of the Markowitz process and what the Markowitz process is or is not. We often forget, in applying financial economics, the basic premise of Markowitz, Miller and Sharpe, and the reasons for which they got their Nobel prizes. What they said was in simple terms and the mathematics followed, not vice versa. If we know what *we* think the expected returns of a set of assets are — not what the market thinks, not what is implied from looking at history — then, given this estimate of expected returns, we need to know the errors of those forecast returns, not what the market has done over the past 15 years, 5 years, 18 months or on average, not whether it is ARCH or any kind of mathematical model. Our model of volatility must be *consistent* with our set of expectations of returns. Then, with the correlations of those errors in our forecasts, we can create optimal allocations. From this simple principle most of the rest of the work in financial economics is derived.

Much of what we do in the profession is creating this information. The basic premise is quite different from the models that we spend most of our time trying to apply to history or to markets themselves. It says that we can create optimum portfolios if we have a set of expectations and if we know what the errors of those expectations are. It says nothing about the models that we use to create those expected returns, just that those three factors of expected returns, the errors of those expected returns and the correlations are all consistently produced.

This is extremely important, because in our industry we spend a great deal of our time on things that I would call 'spurious specificity'. We can model to the *n*th degree, many times and in interesting ways, possible return or risk outcomes, but forget that, in applying those models, we are really looking at questions that are not asked or addressed by the basic Markowitz premise. This can lead to incorrect conclusions. The author does not fall into this trap, but some of the comments that have been made suggest that it is easy so to do.

The basic premise above says something about risk adjusted expected returns in the Markowitz process. There can be nothing else other than have a probability adjusted expected return, after the expected return is determined — to use the terms of Mr Larner — in creating a discounted stream of cash flows or determining a future value for an asset.

It is our estimate of the future and it is our probability of being wrong that is critical; that is, if we were prescient, if we knew the future, *we would not need portfolios nor would we need portfolio theory.* So, in a sense, when we create models based purely upon history, it suggests that we know the future of, say, risk. We are suggesting that we believe that we can be prescient, at least in one area, but portfolio theory exists because we are not prescient in any area. Applying concepts that suggest perfect foresight on a portfolio context makes no sense. This critical, point has been missed by earlier speakers. Our forecasting errors are as critical, or more critical, bits of knowledge, than the actual forecasts of expected returns. These must be consistent with our estimate of a return, or once again we may be making initial mistakes when applying selected propositions.

The author has educated me, as a financial economist, about the actuarial profession, but I want to come back to many of our basic disagreements, which would simply vanish if we would focus on the fundamental premise of financial economics. We do not know the future, and if we do not know the future, then all that we are trying to accomplish is creating a way of getting rid of risks, and those risks are associated with our lack of knowledge of the future.

Professor A. D. Wilkie, F.F.A., F.I.A.: Like other people, I have found this not always an easy paper. Many papers that turn out subsequently to be seminal papers are quite difficult. The author has done a good job in making things fairly intelligible for some of us. I hope that all will re-read the paper.

In ¶2.4.4 the author gives a formula, showing how he derives risk-neutral parameters. If one shifts from *q*s into constant forces of lapse, that I will call μ, and from interest rate *i*s to constant forces of interest, that I will call δ, that formula simplifies into:

$$\mu_{RN} = \mu_{BE} + \delta_{BE} - \delta_q.$$

This is just the same sort of trade-off between μ and δ that we all learnt when studying life contingencies. A constant addition to mortality rates is numerically exactly equivalent to a constant addition to the force of interest. It is convenient in life tables to adjust the force of interest rather than go through the complication of calculating a new life table, but the same thing can be done the other way round. Instead of adjusting the rate of interest, you can adjust the life table.

While there may not be an exact analogy between what the author is talking about in terms of risk neutrality, it gives perhaps more old-fashioned people a feel for what is happening. In some respects it is just a numerical trade-off between different ways of calculating the same thing.

In Section 3 and Appendix A the author discusses dynamic optimisation. I am glad to see this being considered. Quite a lot of work is being done by mathematicians in this country and elsewhere on stochastic dynamic programming methods, which are ways of doing the dynamic optimisation by re-formulating the problem as a linear programming problem. It can be immensely complicated mathematically. As the author has pointed out, dynamic optimisation may take advantage of some of the features of the model. I do not think that it has been noted that this was discussed quite a bit in Wilkie (1986). I showed that, by using my then model and a fairly naive system of looking three years ahead, and using what I would call central estimates (that is ignoring the standard deviations), one would get, possibly, 3% p.a. extra return, which is roughly what the author also found. I pointed out that you could not rely on the market behaving like that. It is unreasonable to assume that one could do this without other people noticing.

Dynamic optimisation is a perfectly legitimate piece of technical investment policy, to try to take advantage of the state of the market at the time, selling shares when they are dear and buying them when they are cheap. You may make a large extra return, but you cannot rely on it, because the model might be wrong. Of course there is a limit to what you can put into any one paper, but the author has not shown what would have happened for each of the models if he had used the rules in Model A and the world behaved like Model B. My guess would be that, if you used the rules applicable to my type of model, and shares actually behave as a random walk, then you waste some expenses, but you do not lose disastrously. If the world is really like my model with different parameters, you might very well lose disastrously. That needs further investigation.

The author has also referred to stochastic models, almost a separate topic. When looking at different models, one should consider what particular elements they simulate or represent. Total return models do not give you dividend yields and interest rates. If you use a completely total return model for your simulation, you do not actually know what the basis is for calculating valuation returns or minimum funding requirement returns. They are incomplete in that sense. The type of model that I have developed may have its problems, but it does give you interest rates and dividend yields as well as total returns.

The author mentions the chaotic model, which is a non-linear deterministic one. If one had the suspicion that it might be an ARCH model, then time series analysis would soon find out. You would investigate the correlation between X_t and X^2_{t-1}, and probably find that there was a perfect correlation between those. You could find out if somebody was using a chaotic model.

Concerning what we call the actuarial value or assessed value for pension fund purposes, I could justify it using my type of model, which is, I hope, a good model of an inefficient market. Let us assume that we know everything that we need to know in order to do a valuation in respect of the assets and the liabilities, except that we do not know anything about market values today. Then we could use what I would call the neutral values or central values from my model for the valuation, or even, which is slightly different, the long-run mean values from my model. It would not be the same as the market value except by chance, but, in a sense, it is what the market value ought to be, not because the market has got it wrong today, but because it gives what the market value on average would be.

The business cycle is discussed in Section 4. The Modigliani & Miller thesis — in fact much of what the author is saying — assumes that tax is not relevant. However, the tax position of the company and shareholders combined is relevant, and British companies are possibly now noticing that rather a lot of their shareholders are their neighbours' pension funds and are gross. If the shareholder is gross and can reclaim ACT, I would not be surprised if the two lines in Figure 4.3.2 overlap, because it seems to me beneficial for the combination of shareholder and company to have maximum dividend distribution and lots of rights issues simultaneously, as the ACT is repaid by the government.

Mr P. D. G. Tompkins, F.I.A.: I confine my remarks to the pensions area, and, in particular, to the comments made in Section 2 about different ways of looking at risk.

I was pleased to see that the author acknowledges, in the concluding section of the paper, that models which seem blatantly unrealistic can sometimes give rise to good management decisions. That is one of the challenges for the theoretician, that the impure work that we carry out does work in many practical circumstances. For example, take the references to best estimates in the paper. If you are faced with the possibility of acquiring a pension fund which has been funded on what might be properly described as 'prudent best estimates', and you wanted to argue for a realistic best estimate of long-term returns, you would be arguing for a lot less than the people selling the business might be prepared to offer.

The author refers, in Section 2, to the actuarial reserve concept, and describes this as the quantity of assets required today to meet future liabilities with a pre-determined probability. It is that pre-determined probability which I find difficulty with. Who is to determine that probability, and how is it to be chosen?

The minimum funding requirements come into force in just over 12 months' time, in 1997, and there has been some difficulty in choosing what level of confidence there should be in the fund being able to meet the liabilities. More illumination from the profession on the mathematics behind this would be particularly helpful, especially as the minimum funding requirement develops and we get experience in its practice.

The author makes reference, in ¶2.2.4, to the use of the different discount rates for different components of distributable profits. This is a matter which was considered at some length in the discussion of Thornton & Wilson (1992). The suggestion put forward there was for a dual valuation approach — the valuation of the accrued liabilities and the assets at the date of the valuation on a different basis from the valuation of the contributions subsequently needed to provide for the continuing accrual of liabilities. The problem here is with communication. How do we impart to the recipients of the advice the understanding that we are assigning a probability to the event that the assets will be sufficient to deliver the promised benefits?

In ¶2.3.7 the paper addresses the fact that scheme valuations are carried out on a best estimate basis. The best estimate he posits is one which will, if returns are then monitored, produce a subsequent return with a mean of the best estimate and random noise. I do not think, in practice, that that is what is meant by 'best estimate', and I think that the term is used glibly.

At the end of Section 2 the author questions the method of determining the appropriate risk adjustment. This is an area where the paper leaves a lot more work to be done. What is the purpose of measurement of risk? Are we concerned with the risk of default? Are we concerned with the risk of the returns not being as high as are expected?

Dr A. J. G. Cairns, F.F.A.: In the example in Section 2 higher lapses mean slightly lower profits. It therefore seems that, in the framework of Section 2.4, risk-neutral lapse rates should be higher than their best estimates. In many other examples it is quite likely that higher lapse rates result in greater profits. In such cases I would expect risk-neutral rates to be lower than their best estimates. If the formula in ¶2.4.4 is applied blindly, then it would follow that we ought to discount earnings at a lower rate than the best estimate. This does not seem to be a particularly rational thing to do. So, I suspect that the risk-neutral approach described in Section 2.4, and the dual risk discount rate approach described in ¶2.2.5, are not consistent, unless my general reasoning has fallen down in some respect.

Risk analysis needs to be thought through carefully. Lapse rates will be subject to the following minimum set of risks: investment returns; other economic factors; randomness in the time to lapse or maturity on the part of each policyholder; errors in the estimates of the basic parameter values (here α and β); and potential errors in the model specification.

Now the third of these (the randomness associated with individual policyholders) is one of the few risks which falls with the size of the portfolio; that is, it is diversifiable. It can, therefore, be argued that this risk can be ignored. However, I believe that this is a simplification, since, for example, mutual companies and smaller investors (for example small pension funds) cannot diversify such risks away entirely.

I have a feeling that deriving risk-neutral parameters will be rather more difficult when such a fuller range of risks is included. However, I would encourage the author and others to develop this approach further, since I see this as one of the major bones of contention between financial economists and actuaries.

One form of assessment was not evident in Section 2, and this was the use of a single risk discount rate. A potential investor may only be interested in net cash flows rather than income and outgo separately. The question then arises as to what the appropriate risk discount rate should be. In the distant past, an actuary may have specified 15%, or perhaps 12% for a mutual. More recently practice has moved in the direction of more rigorous risk analysis, resulting in different rates for different types of contract. In theory, it is possible for us to gauge from the market what risk discount rates are accepted for a given level of risk. Here the definition of the risk discount rate could be that rate of interest at which the expected net present value of the future random cash flows is equal to its market value. If the risk discount rate is higher, then, by implication, the variance of the net present value of the cash flows should be correspondingly higher.

We will then be able to draw a graph which plots risk discount rates against the maximum variance of the net present value. In the example of Section 2 the variance of the net present value will be higher than that for an investment in the underlying stock. Hence a higher risk discount rate will be required. How much higher will depend on the risk that lapse rates and expenses growth deviate from their expected values.

Mr J. P. Ryan, F.I.A.: I shall make a few comments about particular types of models. Actuaries spend most of their time looking at the liabilities side of the balance sheet, utilising observed bases to forecast the way those liabilities emerge. The forecasting models are derived from both a detailed analysis of underlying data and from underlying beliefs of the way that the world works.

Professor Wilford said that, to some extent, that was the way that Markowitz, Miller and Sharpe put together their view of financial economics — basing their model on a view of the world rather than taking a purely market driven approach. Much of this discussion, and a lot of the work in the paper, has centred on market-driven models, derived from a purely observational analysis of market prices, rather than on a fundamental analysis of the way that the world works.

In the paper the author discusses how these two sides can be reconciled, provided the underlying assumptions can be reconciled. I think that the difference between these two fundamental approaches gives rise to some of the puzzles that the author posed in introducing the paper, partly because in asset-liability work there is a great temptation to use: on one side a market-based pricing type of model; and on the other side a more fundamental type of liability model; without necessarily totally reconciling all the underlying variables.

We have seen some movement towards the use of market-based models of liabilities. If you go back over the past 20 to 30 years in the life assurance industry, you can see that risks have been shifted away from the company to the policyholder. One can make an attempt to price some of those risks, as perceived by different companies and by the life assurance industry as a whole.

There are similar developments in the general insurance field, with the introduction of the concept of property futures and property insurance options. This provides a market-based method of actually pricing the risk, rather than the more conventional actuarial approach of building up individual claims.

Similarly, in the Lloyd's market one is seeing the development of traded participations and syndicates which provide another basis for modelling and valuing the liabilities. The purpose of this exercise is not to say either that one approach is right or that another is wrong, but that we can learn something by looking at the two different types of models.

Mr S. J. B. Mehta, F.I.A. (closing the discussion): I was pleased to see, in ¶2.1.2, a robust defence of market value, and I do hope that the profession will increasingly move away from subjective concepts such as long-term value. It was also pleasing to see two specific examples, in relation to the valuation of PEPs and European options, reconciling the use of cash flow specific discount rates and a risk-neutral valuation approach. These showed that the effect of higher investment returns, for example through investment in equities, is offset by the need to use a higher discount rate. This observation has wide implications for actuaries working in both the insurance and pensions fields. For example, it seems clear to me that the methodology proposed by the author in relation to the valuation of pension fund liabilities would not result in the use of discount rates equal to the equity rates of return.

The implications of the author's remarks at the beginning of Section 3 are also far reaching. Investors are able to rearrange the distribution of their total wealth to offset the impact of any departure from apparently optimal insurance company or pension fund asset mix. From this standpoint, it makes no difference, from a shareholder's or member's perspective, whether an insurance company or a pension fund invests 100% in equities or 100% in bonds. Of course, this statement needs to be qualified, but asset/liability studies and asset models need to take account of the implications of investor flexibility.

Section 3 provides a very welcome commentary and analysis of six different stochastic models, including the author's own, and new, jump equilibrium model. This model represents a major advance over many existing models. Indeed, we have heard from the banking world and a derivatives trader some of the ways in which his model may improve on the existing Black-Scholes model and on the work of investment banks. Aside from mathematical elegance and symmetry, it has many new and important properties. Besides realistic means and variances, the model reflects the important characteristic that there are occasional shocks to the market environment.

The model can be used to assess the market price of any derivative or underlying security, of any term, for any of the four main asset classes being modelled. Mean returns are allowed to drift over time with increasing uncertainty over the long run. The term structure aspect is, perhaps, one of its most interesting features, since, for example, the asset class 'sterling' takes as input the existing yield curve on U.K. government bonds. Many of the existing interest rate models do not replicate the yield curve, and therefore lack some credibility when modelling assets or liabilities with more complex characteristics.

The cost of symmetry is a model which permits occasional negative interest rates and shows some drift in real returns. These features may not be of concern in a practical context, but more research into the properties of the model would be desirable.

Further work is needed on the subject of dynamic optimisation and asset strategy. This could supplement the very useful analysis contained in the paper of how asset model characteristics interact with, and affect, the choice of an optimal strategy. In particular, the key question to my mind is, can asset strategy be optimised by taking a year-by-year approach, or are long-term asset return relationships a vital ingredient?

Section 4 introduces a new method of assessing optimal capital for a company, and proceeds to show how analysis of the effects of transaction costs may provide a crucial link to our understanding of the causes of the business cycle. Although developed in a general insurance context, it would be interesting to see how the concepts can be extended to analysis in a life assurance setting.

Section 6 adds considerably to an understanding of how appraisal values can be determined in practice. The apparently arbitrary choice of individual component discount rates in Section 2 is explained. The discount rates are seen to have been selected having regard to systematic risk, as would be expected, given the close analysis and defence of the capital asset pricing model, also contained in the paper.

Appendix A provides an excellent introduction to dynamic optimisation techniques, and gives a good explanation of why utility theory is often more robust than alternative asset/liability optimisation measures. Finally, Appendix B sets out a number of new approaches to the selection of stochastic asset model parameter values.

With the recent discussion of the future of the profession in mind, I believe that inclusion of this paper as required reading in the examination would further improve the status of actuaries within the financial community. The paper is far reaching in its implications for our work, and I would recommend all actuaries, not just the mathematically inclined, to take some time to absorb its contents.

The President (Mr C. D. Daykin, C.B., F.I.A.): What we have heard, fairly consistently, in this discussion is that it is important that actuaries begin to understand financial economics, in spite of the sometimes opaque mathematics, and this paper has helped a great deal in opening up some of the new vistas and possibilities for actuaries to communicate better with financial economists and other people in the world of finance.

We have also heard that there could be scope, through this type of mechanism, for actuaries to be able to contribute more to the field of banking and finance, where actuaries are only involved to a modest extent at the moment.

There is a real debate within the profession as to the relevance of market values and the appropriateness of actuaries determining their own values of quantities which are also determined by the market. My perception is that, over time, the trend is moving more towards a greater predilection for market values. I suspect that this paper will have, to some extent, reinforced that trend.

We have learnt much this evening, and there will be more to be gained from studying this paper and from studying the discussion in some detail when it is published. The author has been very successful in presenting the concepts in a way which has opened up the subject to actuaries, and has emphasised the relationship between actuarial concepts and those of financial economics. We owe a debt of gratitude to the author for presenting a most excellent paper.

Mr A. D. Smith (replying): The opener talked about an actuarial value based on discounted cash flow and consistent with economic value. I fully endorse consistency with economic value, but it does not matter how he calculates it, provided that he gets something which is consistent with economic value. Somebody else could use a different method, and if they get something which is consistent with economic value, that is also fine. Let us concentrate on whether the answer means something, and move away from the details of the method.

Professor Clarkson recommends a framework which is much more complex than that which I have proposed. Does he think that generalising the assumptions will make the mathematics any easier? Of course not; every time that we relax assumptions to a tiny degree, the mathematics takes on a whole new order of magnitude of complexity.

Mr Lee raised some interesting points about my own stochastic asset model, in particular how expected returns behaved in an erratic way over very long time horizons. That certainly gives us something to think about. We need to pay attention to long-term asymptotic behaviour. It is fair to say that my model does not shed very much light on it. One of the reasons for making these models public was that they would be tried out, and the strengths and the weaknesses determined.

Professor Wilkie mentioned dynamic optimisation. If one uses his model to optimise, potentially the market guesses and the optimisation stops working again. What he is saying is that the model would change if we tried to use it, but will not change if we do not.

Both Professor Wilkie and Mr Ryan mentioned the question of valuation in the absence of published market prices. That is a very interesting question — if we cannot buy or sell assets, where does the value come from? However, the world where we can buy and sell assets is hugely different from the world where we cannot. It may not be all that helpful to abstract to a world where you cannot trade, to infer that much about the world where you can.

Dr Cairns made a number of points about Section 2. I intended Section 2 to be about the mechanics of calculation, and to show that various methods were mechanically equivalent. The explanation of where the assumptions came from was in Section 6. As he rightly notes, one could choose other examples where they work the other way round. It was reassuring for me to be able to explain the answers that I got. I suspect that one could come up with examples which I would find rather harder to explain.

The final point that Dr Cairns made is very interesting. It conflicts with the way in which we do much of our assessment of risk. I am proposing that there are risks which shareholders can diversify and risks which shareholders cannot diversify, and that it is important to distinguish between these. Dr Cairns suggested that the discount rate might relate to the individual policyholder lapse risk. I disagree. The consequence of that assumption would be that if you had two insurers which merged, the overall level of risk would appear to go down, and, consequently, value would be spuriously created. That is consistent with risk neutrality, as he pointed out, but I do not think that it makes commercial sense from a shareholder's point of view.

WRITTEN CONTRIBUTIONS

Mr S. J. Green, F.I.A.: I remember two architectural models, one built of cardboard and the other, which was altogether different, being built of stone, wood and steel. When I commented on how realistic it appeared, I was told that it was excellent for wind tunnel tests, but was no use at all for earthquakes, as no one could model the geological strata on which the edifice was to be erected, and, hence, no proper foundations had been built into the model.

I was reminded of this as I read this paper and came across the random walk model, which I cannot believe that anyone still uses to model stockmarkets. Those of us who have spent some time in the real investment world always knew the theory to be fallacious. More than twenty-five years ago mathematicians showed it to be false in the London equity market, in the sense that there was clear statistical evidence of an, albeit weak, correlation between the price of successive bargains. Later similar evidence was found in other financial markets.

The Nobel prize-winning physicist, Murray Gell-Mann, wrote "...many neo-classical economists have preached for years that the price fluctuations in financial markets around the values dictated by market fundamentals constitute a 'random walk', a stochastic process.... But it has now been shown convincingly that the chance process idea is wrong. These fluctuations are, in fact, pseudorandom, as in deterministic chaos; in principle they contain enough irregularities for one to make money out of them."

Similarly, although the fallacious assumptions upon which Black-Scholes has been based have been well known to the investment community since the model was first produced, the fact that it was universally used by market-makers who could, and did, make money using it has permitted it to become an accepted nostrum and widely used by investors who, because of expenses, could only, coincidentally, profit by its use. Control systems which have been built on this model have not prevented disasters.

To quote from two other physicists, Bouchaud & Sornette in *Physics Today*, March 1996, "The usual perception among traders is that the risks involved in derivative trading must be small in practice because the risks can be made to vanish in theory. In fact no banker using the Black-Scholes equation can know exactly the exposure of his establishment to such risks. This is because the Black-Scholes theory is silent on the crucial question of risk in that the underlying hypotheses and mathematical framework lead to a solution without risk. One cannot measure something which does not exist."

Compared with the crude cardboard boxes which represented the earlier attempts by financial economists to model the financial markets, the author's preferred model may have reached the stone and steel model stage that I described earlier. He has certainly recognised that markets and economies behave in a more complex manner than do other model makers, particularly the asset-liability model makers, whose published work seems so crude. Even he, however, has to make assumptions which he knows to be unrealistic — for example in ¶5.3.4 — where he assumes that "bonds trade frictionlessly"; but at least he recognises that "the real world is much more complex than the models which we are trying to fit".

It is interesting that the author describes a model which he uses for companies in the non-life insurance industry. Stripped of its modern mathematics, and assuming that the author pays rather more attention to the size of the market and the number and resources of the other players than he, of necessity, describes in the paper, he is reproducing the sort of fundamental analysis which has now been dropped from the syllabus, even in its most elementary form. A prime example of this was the work which Plymen carried out on the shoe industry for a manufacturer some years before the author was conceived. Although this was based on the derided actuarial tools of probability and compound interest, it provided a reasonable scenario and accurate forecasts.

The models which the author propounds may well be an advance on those produced by the old-fashioned techniques, but only time will tell. In the meantime, may I caution all members of the profession to apply their common sense, as well as their new models, by recounting the story of the leading consultant advising one of my former clients. The pension scheme consisted of several thousand pensioners and deferred pensioners, for whom index-linked annuities had been purchased from a well-known life office, together with a hundred or so active members entitled, like the pensioners, to index-linked pensions. Of course index-linked annuities do not, usually, replicate index-linked pensions, as the former may increase or diminish whereas the latter can only increase. There is, therefore, a contingent liability in the event of negative inflation. Because of the numbers involved, this contingent liability far exceeded any others in the scheme. Fortunately there was sufficient surplus in the fund to cover this liability. The consultant, using one of the most sophisticated asset-liability models produced by financial economics, advised the client to invest the surplus in index-linked securities, and is still to be convinced that his advice was illogical.

Mr J. M. Pemberton, F.I.A.: The title of this paper suggests that it addresses the most important issue facing the actuarial profession today — "what is the role of financial economics in relation to actuarial science?" However, the first six sections discuss a number of techniques used within economics or actuarial science, but fail to identify whether each of the methods is actuarial or part of financial economics. A careful study of the text book development of financial economics will reveal that the subject has at its core a disregard for empirical investigation — it has no tradition of checking its claims by reference to the facts. Such an investigation reveals that there are important fundamental differences between financial economics and actuarial science — the author's attempt to create a synthesis between the two by ignoring the differences is misguided.

Specific criticism of individual errors is beside the point — the paper fails to grasp the real issues of the important underlying debate, because it is trapped in a discussion within the subjects (financial economics and actuarial science) rather than engaging in a discussion about these subjects. Indeed the author, like those who engaged in the discussion, seems to be unaware of the existence of the huge body of literature about economic methods which is relevant to our debate. The paper, 'The methodology of positive economics' by Milton Friedman, from *Essays on Positive Economics* (University of Chicago Press, 1953), heavily influenced by the work of Popper, represents a landmark development within this area. However, there have been two revolutions within the study of method in science since then. Amongst the hundreds of other intellectual heavy-weights who have contributed to this debate are Ernst Nagel, Thomas Kuhn, Daniel Hausman and Nancy Cartwright. Mathematical games are futile unless we know how to use the resulting models to tell us something about reality, although the author does briefly acknowledge this debate, in Section 7, in discussing the relationship between models and reality.

There are many areas in which the tension between financial economics and actuarial science is apparent — in practice the two approaches often lead to very different answers. In my view such discrepancies are worthy of serious investigation. If we wish to find a synthesis of the two approaches, the first place to start looking is in areas where they provide apparently conflicting answers. I believe that the discrepancies in these areas are due to deep-seated and fundamental incompatibilities within the two sets of methods. If this is right, we must make a choice between the two approaches.

My investigation of the relationship between models and reality within financial economics leads me to be very cautious about the value of such models. I believe that many models in financial economics often provide a good approximation to reality, but these same models are sometimes very wrong. I am concerned about the acceptance of financial economic models by actuaries, particularly where this implies a choice against actuarial models, without a sufficient understanding of the nature of such models.

I am deeply disappointed by the way, on the evidence of this debate, that the profession has been seduced wholesale by the excitement of financial economics. Many of those who are at the forefront of the Institute work in this area complain that the mainstream of our profession is 20 years behind the game. As a result of my investigation of financial economics from a methodological standpoint, I have no doubt that financial economics has serious limitations. I believe that the actuarial profession has precisely the tools necessary to provide a scientific basis for addressing the relevant questions in new and successful ways.

The author subsequently wrote: A number of speakers asked about appropriate discount rates for measuring the economic value of pensions liabilities. The valuation problem is immediately complicated by the number of other judgemental assumptions also involved — for example the rate of salary inflation. In addition, as a number of speakers have pointed out, when we talk about expected cash flows we may not be using 'expected' in its strict statistical sense. However, I agree with Mr Exley on the fundamental point that, at least for a creditworthy sponsor, the economic cost of a pension promise should not be directly affected by the assets that the scheme happens to invest in.

The general economic principle in Section 6 is that the economic value of a liability is the market value of a matching asset portfolio, that is, the minimum risk position. The opener has suggested, instead, that one might consider "an investment policy which is appropriate to the nature and risks of the liabilities", which does not necessarily mean minimum risk. Many funds take calculated mismatch risks in the hope of securing better returns, so, if a fund discounted liabilities using the expected returns on the assets actually held, they would be taking credit for these returns in advance of their having been achieved. However, if the asset-liability match is not close, then the sponsor's capital base is being put at risk, so, in shareholder value terms, the need for a cost of capital adjustment would undo the effect of a high equity-related discount rate.

Matching considerations can be rather more subtle than would appear at first sight. In the long run, perhaps, corporate profit tends to be carved up in fixed proportions between capital and labour, so that there may be a broad match between dividend and wage inflation, and hence between equities and pensions liabilities. The limitation of this argument is that a particular scheme will not receive the aggregate dividends; it will miss out on those arising from new issues. Furthermore, the scheme will not pay aggregate salary inflation, even pre-retirement — it will pay the specific rises awarded to actives, and another rate to deferreds. One could argue that, in an era of corporate restructuring, the asset decrement of new issues may be offset by higher decrement rates on the liability side from active to deferred, perhaps because of a redundancy programme, so that equities remain broadly a good match for liabilities in respect of active members. I have not examined these aspects numerically, but it does seem important to consider, not only the relative sizes of aggregate cakes, but also a more detailed assessment of the proportion that applies to a particular fund. I remain to be convinced that 100% of the equity risk which justifies its return premium can also be applied to the liability side. For example, the salary of an individual cannot, generally, decrease from one year to the next. No such promise is made for equity dividends. My understanding of the current MFR proposals is that, in the short run, certain long-term liabilities are matched by equities 'by construction' rather than because the relevant issues have been properly thought through.

The issue of default risk complicates matters. Mr Lee has pointed out that the capital of the sponsor is not a total guarantee for the pensions promise, and so the cost of the capital actually provided may be less than my calculations suggest. This introduces a very important point, that the economic value of future pensions depends, to some degree, on the creditworthiness of the sponsor and the funding level of the scheme. A decision by the scheme trustees to speculate in equity markets may well reduce the economic cost of providing the pensions by increasing the probability that the pensions promise will not be fully honoured. An alternative view is that speculative investment strategies are bad for the sponsor, because trustees may press for out-performance to be spent on discretionary increases. However, the essential point remains that the trading position of the sponsor and scheme ought to have a significant impact on pension valuations.

Fortunately, although we may claim to apply expected equity returns to discount pensions liabilities, in reality valuation bases rarely reflect what I would regard as 'realistic' equity returns by reference to historic data. As Mr Tompkins has pointed out, the answer may make commercial sense, even if the premises are indefensible. Some actuaries assume a modest risk premium of, say, 1% for the expected return on equities versus gilts when valuing pensions. In a forthcoming paper, Mr Exley and Mr Mehta argue that such a figure for the liability discount rate may be economically valid, not for the stated reasons, but on the grounds of credit risk or because of the correlation between redundancies and the economic cycle. 'Equity returns' have, in effect, become a euphemism for savings in pension cost arising from redundancy or default. It would not inspire the confidence of scheme members to learn the true reasons for the choice of liability discount rate. Furthermore, if we are interested in assessing the relative economic costs in respect of active, deferred or retired members, it seems to me that different discount rates would be appropriate in each case.

It is good that a wide range of views were represented, particularly in the written contributions. It is clearly sensible to scrutinise our models thoroughly and from a suitably sceptical perspective.

Mr Pemberton would have us believe that financial economics and actuarial science are two completely independent schools of thought competing to solve the same problems, the former school being flawed, but the latter sound. However, I am hard pressed to identify a purely actuarial theory of derivative pricing, or of shareholder value, or of dynamic optimisation, or of asset-liability management. Quite rightly, actuaries have been active in these areas, adding our own techniques to the melting pot, but we are all deeply indebted to financial economists for much of the methodology in practical use.

Mr Green cannot believe that anybody still uses the random walk to model stock markets. Certainly, such models are of limited use for active fund managers. When performing asset-liability studies it is conventional to take a broad view of the assets, assuming, perhaps, that they track some index. There is some loss of detail here — if the manager is actually picking stocks on the basis of some technical method involving, for example, Fibonacci series, then he may not track the index. He may hope to do better, but may actually do worse. I am not aware of any asset-liability methodology currently in use which takes detailed account of the processes actually being employed by the fund manager. However, such alternative processes do not seem to be assisting managers in their desire to outperform. Further, I do believe that simple random walk models do have a role to play in many circumstances where resources are scarce.

Press reports have exaggerated my differences with Mr Clarkson. There is always room for improvement in current models. Mr Clarkson's verdict is that we are using the 'wrong kind of mathematics', but I would suggest that it is not only quality which is lacking, but quantity. Generalising assumptions invariably involves an increase in the complexity of the mathematics. If we are not prepared to accept 'standard' assumptions, then we are implicitly suggesting that the maths is too simple. In the short term, limited mathematical skills force me to accept assumptions which, on the face of it, are untenable, in order to get numerical output. Over the longer term, I hope to learn some more mathematics which, I am confident, will enable me to build better models. I may even end up somewhere near where Mr Clarkson would like us to be. A pragmatic approach is in order.

One thread running through many of the contributions is the concept of value. Should we have a concept of 'actuarial value' which distinguishes us from other professions? I was recently questioned at a derivatives conference on how options might be valued in the context of a pension fund, using actuarial values. This is particularly topical as funds examine derivative solutions to protecting MFR solvency. It soon became clear that the audience could not agree on the more fundamental issue of the treatment of new cash. Overseas equities are also problematic. What we have at the moment is not a proud professional tradition, but a hopeless muddle. Ambiguity on our part has deterred funds from implementing derivative solutions which could provide substantial commercial benefits. Dyson & Exley (1995) described a market-based approach, which seems to me to be much more satisfactory. At that time, some of the contributors to the debate were not convinced. However, on the strength of this debate, I believe that the door now stands ajar for a thorough review. We should not let the opportunity pass.

REFERENCE

EXLEY, C.J. & MEHTA, S.J.B. (1996). Asset strategy for defined benefit pension schemes. Paper presented to the Institute & Faculty of Actuaries Investment Conference.