DISCUSSION PAPER PI-0102

Pensionmetrics: Stochastic Pension Plan Design and Value-at-Risk During the Accumulation Phase

David Blake, Andrew Cairns and Kevin Dowd

October 2001

ISSN 1367-580X

The Pensions Institute
Cass Business School
City University
106 Bunhill Row London
EC1Y 8TZ
UNITED KINGDOM

http://www.pensions-institute.org/

ELSEVIER

Insurance: Mathematics and Economics 29 (2001) 187–215

INSURANCE MATHEMATICS & ECONOMICS

www.elsevier.com/locate/econbase

Pensionmetrics: stochastic pension plan design and value-at-risk during the accumulation phase

David Blake [a,*], Andrew J.G. Cairns [b], Kevin Dowd [c]

[a] *Pensions Institute, Birkbeck College, University of London, 7-15 Gresse Street, London W1T 1LL, UK*
[b] *Department of Actuarial Mathematics and Statistics, Heriot-Watt University, Edinburgh EH14 4AS, UK*
[c] *Nottingham University Business School, Jubilee Campus, Nottingham NG8 1BB, UK*

Received 1 January 2001; received in revised form 1 May 2001; accepted 1 June 2001

Abstract

We estimate values-at-risk (VaR) in the accumulation phase of defined-contribution pension plans. We examine a range of asset-return models (including stationary moments, regime-switching and fundamentals models) and a range of asset-allocation strategies (both static and with simple dynamic forms, such as lifestyle, threshold and constant proportion portfolio insurance).

We draw four conclusions from our investigations. First, we find that defined-contribution (DC) plans can be extremely risky relative to a defined-benefit (DB) benchmark (far more so than most pension plan professionals would be likely to admit). Second, we find that the VaR estimates are very sensitive to the choice of asset-allocation strategy. The VaR estimates are also sensitive, but to a lesser extent, to both the asset-returns model used and its parameterisation. The choice of asset-returns model is found to be the least significant of the three. Third, a static asset-allocation strategy with a high equity weighting delivers substantially better results than any of the dynamic strategies investigated over the long term (40 years) of the sample policy. This is important given that lifestyle strategies are the cornerstone of many DC plans. Fourth, conservative bond-based asset-allocation strategies require substantially higher contribution rates than more risky equity-based strategies if the same retirement pension is to be achieved. © 2001 Elsevier Science B.V. All rights reserved.

JEL classification: C15; D81; D91; E27; G11; G23

Keywords: Defined-contribution pension plan; Value-at-risk; Model risk; Asset-allocation strategy

1. Introduction

The world of pension provision is currently undergoing two dramatic transitions — a shift from unfunded social security towards private funding and, within the private funded sector, a shift from defined-benefit (DB) plans towards defined-contribution (DC) plans. Whatever their merits, these shifts involve enormous transfers of risks from taxpayers and corporate DB sponsors to the individual members of DC plans.

The question then naturally arises: how successful will DC plans be in delivering good pensions in retirement? The blunt answer is that no one knows: not the DC plan provider who has no special obligation to the plan

* Corresponding author. Tel.: +44-20-7631-6410; fax: +44-20-7631-6416.
E-mail address: d.blake@bbk.ac.uk (D. Blake).

PII: S0167-6687(01)00082-8

member if the terminal fund value is low; not the government, which may come under pressure to bail out members if realised pensions turn out to be very low; and certainly not the plan member himself who is being sold the *promise* of a comfortable retirement on the basis of modest contributions and high *anticipated* investment returns. There is a certain irony here. Car manufacturers spend many millions of dollars in designing each new model. The reason is clear: the potential customer can tell immediately whether the product is any good. But in the case of a pension plan, which might last 70 years or more before its cycle is finally completed, the consumer will not know whether the product is any good until he is on his deathbed and then, of course, it is too late. That we believe is why we do not observe similar sums being spent on the design of pension plans. Evaluating pension plan design was always important, but is becoming even more so as the move towards DC plans gathers pace.

A good DB plan typically offers a pension that is related to the member's final salary achieved just prior to retirement and is subsequently indexed to inflation until death. The financial risks associated with DB plans are borne by the plan sponsor, usually a large company, rather than the plan member. However, most DB plans suffer from poor portability, so that when a worker moves jobs, he can end up with a much lower pension in retirement. [1] To this extent, some risks are borne by the worker, namely the economic risk that the pension from a previous employer revalues at a lower (and uncertain) rate relative to the worker's salary, and the demographic risk from the number and timing of job changes and periods of unemployment. The low portability of DB benefits creates a major disincentive against moving jobs, and is therefore a considerable hindrance to labour market flexibility.

By contrast, a DC plan does not suffer from portability problems, since it is easily transferable between jobs. However, DC plans impose a range of risks and other problems on the plan member. He may have to pay excessive charges and other costs imposed by the DC plan provider (e.g., for marketing, administration, fund management, etc.). He bears the risk that inadequate contributions might be made into his plan (e.g., due to spells of unemployment, child care or ill-health or simply because the contracted contribution rate was not high enough). He also bears asset price risk (the risk of losses in the value of his pension fund due to falls in asset values, especially in the period just prior to retirement). As he retires, he bears interest-rate risk (the risk of retiring when interest rates are low, so that the retirement annuity is permanently low). After he retires, he bears, if he purchases a level annuity, inflation risk (the risk of losses in the real value of his pension due to subsequent unanticipated inflation) and, if he purchases an investment-linked annuity, income risk (the risk that the payments made on the annuity will fluctuate). Throughout, he bears the risks of unfavourable changes in regulatory regimes, e.g., the imposition of onerous pension fund investment restrictions. Further, the insurance companies selling annuities face reinvestment risk (the risks associated with a failure to match asset cashflows with expected liability outgo) and mortality risk (the risk that their pool of annuitants has lighter mortality than allowed for, e.g., because of an underestimate of mortality improvements).

This paper sets out a simple and practical-to-implement methodology to evaluate alternative pension plans. In principle, this methodology enables us to evaluate and compare any plan using any yardstick chosen. Here we use the methodology to assess how closely DC plans can replicate the benefits available from DB plans. In particular, we aim to design a DC plan that can target a particular pension outcome (or pension relative to final salary) with a specified degree of probability, and in so doing quantify the degree of pension risk — the potential difference between the DC pension and a target final-salary pension — inherent in a DC plan. We do so by estimating the values-at-risk (VaR) of the pension funds implied by alternative DC schemes and then use these VaR estimates to make a formal comparison between the two types of scheme. [2]

[1] To give an illustration, Blake and Orszag (1997) found that a typical UK worker moving jobs six times in a career could end up with a pension of only 71–75% that of a worker with the same salary experience who remains in the same job for his whole career.

[2] The assessment and management of investment and contribution risk in a DB plan needs, primarily, to be considered from a plan sponsor's point of view. This is dealt with by Cairns (2000a).

VaR is now a well-established risk measure, and has been applied successfully in other fields such as banking and risk management (Dowd, 1998).[3] In using VaR here, we need to take a view on the confidence level(s) used to make the comparison. Suppose we wish to achieve the same replacement ratio[4] as that generated by a specified DB plan. Our simulations will produce an empirical distribution of possible pension ratios[5] for our DC plan, and the range of values of this distribution will typically vary from substantially less than unity at the lower end to well above unity at the upper end. To make our comparison, we need to specify one or more percentiles from our distribution, and then compare these values with our target pension ratio of unity. We would take the *i*th percentile, which is the VaR at the $(100 - i)$th percent confidence level. If this percentile is greater than or equal to unity, we would conclude that the DC plan successfully replicates the DB one; if the percentile is less than unity, we would conclude that the DC plan does not successfully replicate the DB one. For the purposes of illustration, we made our comparisons with three different percentiles or VaR-confidence levels: the 50th, 20th and 5th percentiles, corresponding to the VaR at the 50, 80 and 95% confidence levels.[6]

This simulation methodology has a number of attractive features:

- The model is extremely flexible and can accommodate almost any set of assumptions or features relating to existing types of pension arrangements. The model therefore has considerable practical as well as academic potential.
- The methodology allows us to carry out 'what if?' experiments and stress tests by changing key assumptions and observing how these changes affect our results. These exercises are very useful because they enable us to identify the driving forces behind our results and gauge the sensitivity of our results to particular assumptions.
- The model can be extended beyond the accumulation phase (the period up to retirement) to deal with the distribution (or post-retirement) phase as well. We report on this extension in a companion paper (Blake et al., 2000b).

The layout of this paper is as follows. We begin in Section 2 with a discussion of the methodology of stochastic pension plan design and what it can achieve. Section 3 then explains the accumulation model — its structure, underlying assumptions and potential extensions — and Section 4 presents and discusses some simulation results for the model. Our conclusions are presented in Section 5.

2. Stochastic pension plan design

The appropriate vehicle for designing and stress testing any pension plan is stochastic simulation. This approach generates a range of outcomes (that is, a probability distribution function) for the value of the pension fund at any

[3] VaR has proven itself to be immensely useful as a measure of financial risk. The VaR is also a very natural risk measure for our purposes, where we are concerned with the distribution of the ratio of pension to final salary. However, the VaR is not the only risk measure we could use, and a good alternative is expected shortfall, as advocated by Artzner et al. (1997, 1999). The expected shortfall risk measure has the advantage of being always subadditive — it guarantees that the combined risk of two or more positions is never more than the sum of the individual risks — whilst the VaR sometimes fails to be subadditive. However, we would argue that the VaR is simpler and more natural for our purposes, and, in any case, the VaR is often subadditive anyway (e.g., when risks are elliptically distributed). We would also argue that the subadditivity issue is more relevant in banking applications, where the risk measures are calculated separately for a number of different lines of business and then aggregated, and rather less relevant in the present pensions problem where we are calculating a single VaR for the total risk for an individual policyholder.

[4] The ratio of initial pension to final-year salary. In the UK, a typical DB plan pays a pension equal to two-thirds of final salary to a member with a full service record.

[5] Ratios of DC pension to DB pension.

[6] Greater experimentation is needed to determine the most appropriate percentiles/confidence levels to use. The lower the percentile or the greater the confidence level, the more demanding the comparison and the greater the chance of the DC model failing the test. It would presumably make sense for the confidence level to be at least 50% (so the median pension ratio is at least as great as the desired level for the DC scheme to pass). But it also makes sense to add some margin to cover any shortfall (that is, so we would want a VaR confidence level of over 50%), but we do not want too great a margin (or confidence level) since many otherwise good DC schemes will fail the test.

given future date, conditional on various assumptions about pension fund contributions, asset returns, mortality, and other relevant factors. This approach involves the following steps:

1. We begin with the parameters defining both the DC pension plan and the stochastic processes driving asset returns, earnings growth, unemployment, mortality, etc.
2. We choose an investment strategy (a set of contribution rates and an asset-allocation strategy) for the relevant investment horizon. Contribution rates may be constant, age-dependent or dependent on the value of the DC assets; similarly, the asset-allocation strategy may be static (that is, 'buy and hold') or dynamic (that is, involve periodic portfolio rebalancing); in the latter case, the asset-allocation strategy may be backward-looking (e.g., the portfolio may be rebalanced in response to realised gains or losses) or forward-looking (e.g., the outcome of a stochastic dynamic programming exercise).
3. Given these various assumptions, the program will generate an empirical distribution of possible pension ratios at any specified future date, a key date being the nominated retirement date of the plan member. This distribution in turn gives us VaR estimates for any chosen confidence level.
4. If the VaR estimates are unsatisfactory, we increase the contribution rate to an appropriate level (holding the asset-allocation constant). This is done by making reference to the relevant percentiles of the existing empirical distribution.
5. At this point, we have a DC plan that adequately matches the DB plan by our selected standard of comparison, and we can then consider the relative costs of the two schemes.

No DC pension plan can achieve a target pension outcome (e.g., the same pension as a DB plan) with 100% probability; however, it is possible to use stochastic simulation to determine the contribution rate, asset-allocation strategies and other factors (e.g., the contingency reserves required for possible mortality improvements) to achieve the target at a chosen degree of probability (such as 50, 80 or 95%). Naturally, the higher the required probability, the higher the contribution rate, the more conservative the investment strategy, the higher the necessary contingency reserves, and so on.

3. The accumulation model

We will illustrate the accumulation model using the simple case of a typical UK male employee who starts a DC pension plan on his 25th birthday, lives at least until his 65th birthday [7] when he retires, and makes contributions to his pension plan throughout his working life; he has no dependants. He faces asset-return risk, interest-rate risk and unemployment risk. Asset-return risk affects the value of his pension fund, given his contributions; interest-rate risk affects the value of his pension when he retires, through its impact on the value of the annuity from which his pension is paid; and unemployment risk affects his income and, hence, his ability to make contributions to his pension fund.

We have chosen to work with annual returns in real terms. We must also make assumptions about the risk factors and the various control variables and assign values to the parameters of the model.

3.1. Risk factors

3.1.1. Asset returns

The first risk factor relates to asset returns. A body of evidence appears to support two important hypotheses concerning the returns on long-term assets such as equities, bonds and property. The first is that they exhibit

[7] We assume that if the plan member dies before his nominated retirement date, the accumulated fund passes to his estate.

some degree of long-term mean reversion.[8] The second is that there exists a positive long-term equity risk premium: the long-term return on equities exceeds that on ‘safer’ assets, such as bonds.[9] These properties are crucial for fund managers because they help to justify the fund management strategy known as ‘lifestyling’ or ‘age phasing’. Samuelson (1989, 1991, 1992)[10] argues that when asset returns are mean-reverting, it is rational for long-horizon investors, such as pension funds, to invest more heavily in ‘high-risk’ equities than in ‘low-risk’ bonds during the early years of a pension plan (thereby benefiting from the equity risk premium) and then to switch into bonds as the horizon shortens. Indeed, it can be argued that over a long investment horizon where pension plan liabilities are linked to the long-term growth rate of the economy, it is fixed-income bonds rather than equities that are the genuine high-risk asset (although this depends very much on the precise measure of risk being used).

A further justification for a heavy investment in equities in pension plan portfolios rests on the long-term stability of factor shares in national income and on the requirement of finding suitable matching assets for the accruing pension liabilities. Since there are no financial assets in existence that are linked directly to the growth rate in labour income (because no government or, indeed, corporation has yet issued wage-indexed bonds), pension fund managers must look elsewhere for matching assets. If factor shares do not trend significantly over time, equities (which reflect the return on capital) and property become natural long-term matching assets for these liabilities. Nevertheless, the short-term correlation between equity returns and earnings growth is very low and so any short-term justification for equities will rely more on their high relative returns than on their contribution to reducing portfolio risk. However, wage increases are known to be highly correlated with price inflation:[11] this suggests that index-linked bonds should provide a lower-risk, although also lower-return, match for earnings-linked liabilities than equities.

If it turns out that asset returns are generated by pure random walks, the optimal asset allocation may (for certain objective functions) not depend on the length of the investment horizon and time diversification will not be effective.[12] It may also be the case that the size of the equity risk premium is not very precisely determined, even in economies with long and continuous histories of financial market data. Recent evidence further suggests that the equity risk premium may be fairly low in most of the world’s financial markets.[13]

There is another important property of equity returns that fund managers must take into account, namely leptokurtosis. The fat tails in the distribution of equity returns raises the likelihood of extreme outcomes, both positive and negative. Over the long life of a pension plan, a number of large positive and negative shocks to equity returns may be experienced. Of particular concern to DC pension plan members will be large negative shocks. But, as Miles and Timmermann (1999) point out, the timing of the shocks is also important. If a large negative shock occurs just after a plan is implemented, this will have a negligible impact on the fund value, but if the shock occurs just prior to retirement, the outcome could be disastrous. Lifestyling is, of course, designed to reduce the impact of such an event, but the need to implement a lifestyle strategy depends to some extent on the nature of the leptokurtosis.

We assume that the pension fund is invested across six risky assets: UK equities, UK bonds, UK property, UK T-bills, US equities and US bonds. The first four are the key domestic asset categories facing any UK

[8] For example, Lo and MacKinlay (1988) found that the weekly return on US shares are positively correlated over short periods, while Poterba and Summers (1988) found the returns on US shares were negatively correlated over long periods. Similar evidence of long-term mean reversion for UK assets is found in Blake (1996).

[9] Siegel (1997) shows that US equities generated higher average returns than US treasury bonds and bills in 97% of all 30-year investment horizons since 1802. CSFB (2000) shows that similar results hold for the UK.

[10] See also Van Eaton and Conover (1998).

[11] Thornton and Wilson (1992).

[12] Samuelson (1963) and Merton and Samuelson (1974). This is the case, in particular, if the investor has a power or logarithmic terminal utility function.

[13] Brown et al.’s (1995) analysis of all the world’s stock markets that have been in operation since 1900 indicates that globally equities outperformed bonds by a negligible amount on average. Goetzmann and Jorion (1997) argue that the high real return of 5.98% on US equities between 1921 and 1995 can be explained by the high level of political stability in the US, as evidenced by an absence of revolution, nationalisation or financial collapse. The average equity return in other parts of the world was just over half that of the US at 3.29%.

pension fund manager, and the latter two assets can be regarded as proxies for the key international assets they face.[14] Appendix A, Table 5 confirms the above stylised facts about real asset returns: namely that long-term assets appear to exhibit mean reversion, that there appears to be a positive long-run equity risk premium, that most assets exhibit leptokurtosis, and that the contemporaneous correlation between financial asset returns and real earnings growth is not strong. We also find evidence that the real yield on T-bills exhibits strong persistence over time: however, on economic grounds we rule out the possibility that the real yield on T-bills is non-stationary, since this implies a positive probability of the yield wandering off to $+\infty$ or drifting down permanently to 0. A good model of asset returns must account for the statistical properties identified in Appendix A, Table 5.

Knowledge of the true model generating asset total returns is, of course, not possible. However, we have experimented with a number of well-known asset-return models:[15]

- Stationary moments models:
  - multivariate normal model,
  - mixed multivariate normal model,
  - multivariate *t* model,
  - multivariate non-central *t* (NCT) model,
  - bootstrap model;
- Regime-switching model: Markov switching model;
- Fundamentals model: Wilkie (1995) model.

Further details on these models are given in Appendix B. With the exception of the Wilkie model, they are parameterised in a way which is consistent with the data discussed in Appendix A. The Wilkie model parameters, listed in Wilkie (1995), are based on a longer run of historical data for a range of financial indicators (e.g., retail price inflation, share dividend yields, bond yields and foreign exchange rates) which are different from those used in the present analysis.

3.1.2. Interest rates

We need to model the evolution of interest rates over time in order to establish the prevailing annuity factor at the member's retirement date. This is the present value, at the time of retirement, of one unit of an annuity for the remaining life of the annuitant. We divide the pension fund on the retirement date by the annuity factor to derive the pension annuity. The discount rate used to derive the annuity factor will typically equal the yield on long-term bonds. We therefore need a model of long-term bond yields that is consistent with the above asset-return models. With the exception of the Wilkie model, all of the models considered here model total returns on the assets only. It is necessary, therefore, to derive (at least approximately) relevant interest rates in a way which is consistent with returns on cash and long-term bonds.

To achieve this, we propose the following model. Suppose that the gross real return on long-term bonds between times $t-1$ and t is given by $R_{Bt} = (1 + r_{Bt})$, while the gross real return on T-bills is given by $R_{ft} = (1 + r_{ft})$. The achieved, logarithmic risk premium on bonds is defined as $\pi_{Bt} = \ln R_{Bt} - \ln R_{ft}$.

[14] We take US equities and bonds as proxies for international equities and bonds because consistent alternative international return series do not appear to exist for our data period (that is, the last 50 years). In any case, taking the US to be a proxy for the rest of the world is not a particularly unreasonable assumption in the context of this paper. Despite their obvious desirability we do not include index-linked bonds in this analysis. This is because their short history (they have only been in existence in the UK since 1981) did not allow us to generate reliable estimates of their distributional properties.

[15] The list is not exhaustive and there are other specifications we could have used to model returns, e.g., exponentially weighted moving average approaches (e.g., as in Morgan, 1997), and an arbitrage-free model (e.g., Flesaker and Hughston, 1996; Rutkowski, 1997; or Cairns, 1999a).

Suppose that the nominal yield on bonds (y_{Bt}) is determined by the following equation:[16]

$$\ln y_{Bt} = \overline{\ln y_B} + f_1(\ln y_{Bt-1} - \overline{\ln y_B}) + f_2\pi_{Bt}, \tag{3.1}$$

$$\ln y_{Bt} = (1 - f_1)\overline{\ln y_B} + f_1 \ln y_{Bt-1} + f_2\pi_{Bt}, \tag{3.2}$$

$$\ln y_{Bt} = f_0 + f_1 \ln y_{Bt-1} + f_2\pi_{Bt}, \tag{3.3}$$

where $\overline{\ln y_B}$ is the mean of the log of the yield on long-term bonds. Eq. (3.1) incorporates long-term mean reversion as well as a short-term responsiveness to the excess return on bonds. If bond prices rise by more than anticipated then bond yields will fall by more than anticipated: we therefore expect to find $f_2 < 0$. Using ordinary least-squares, we derived the estimates $f_1 = 0.90$ and $f_2 = -1.097$ by regressing historical values of $\ln y_{Bt}$ on $\ln y_{Bt-1}$ and π_{Bt}.

We need to simulate the bond yield for the time, T, in which the pension plan member retires at age 65. This is necessary to derive the annuity factor. Suppose this yield is denoted y_{BT}, then the annuity factor is given by

$$\ddot{a}_{65}(T) = \sum_{k=0}^{\infty}(1 + y_{BT})^{-k}\,\mathrm{Pr}(\text{survive from age 65 to age } 65 + k).$$

3.1.3. Earnings

While the starting salary affects the monetary value of the retirement pension fund and therefore the plan member's pension, it does not affect our key variable of interest, namely the pension ratio. We do however need to make assumptions concerning both the particular life-time earnings profile of the member[17] and the distribution of real earnings growth.[18]

Life-time earnings profiles are typically concave in shape to reflect the fact that the greatest earnings growth usually takes place early in the career (as a consequence of rapid early promotion or frequent job changes, both of which are associated with the receipt of substantial earnings increments), while earnings growth tends to slow down from mid-life onwards, as individual careers stabilise and workers come to rely more on cost-of-living increases than experience and promotional awards.

In some industries (e.g., coal mining), life-time earnings reach a peak well before retirement age (e.g., as miners move from coal-face to surface jobs). Similarly, in some professions (e.g., managerial), the highest paid members often take early retirement[19] or are made redundant in their late 50s and early 60s. This reduces median earnings within each cohort between ages 50 and 65. However, sometimes the apparent peaking of wage profiles is the result of an inter-cohort effect: e.g., managers in their 50s could be more productive and therefore have higher earnings than managers in their 60s because they are better educated and more of them went to university, yet information on both cohorts is used to construct the wage profile.

It is important to model wage profiles accurately, since their shape will influence both the timing of contributions into the plan and forecasts of final salary. Wage profiles for men and women in different professions in the UK can be found in the New Earnings Surveys (2001) (or other sources such as OECD studies). The wage profile used in this paper belongs to a 'typical' UK worker as reported in Adams (1999). It is shown in Fig. 1 with median earnings at different ages expressed as a ratio of average industry earnings. Given the preceeding discussion, it is clear that the

[16] The justification for this structure lies in its consistency with the Vasicek (1977) model, which has a similar relationship between bond returns and bond yields.

[17] This is the earnings-age schedule that determines how the member's earnings change relative to average earnings as he progressively ages. They are sometimes also called wage profiles or salary scales.

[18] We also need to make a subsidiary assumption about the impact, if any, of unemployment on subsequent earnings. We make the simple assumption that unemployment has no effect on subsequent earnings, so when the individual is re-employed after a spell of unemployment, he receives the same earnings he would have done had he worked all along.

[19] Such individuals are likely to have built up above-average pension entitlements for their age group both on account of their high salaries (which gives them a high pension fraction) and their long service (which gives them a high pension multiple and is implied by their above-average salaries for their age).

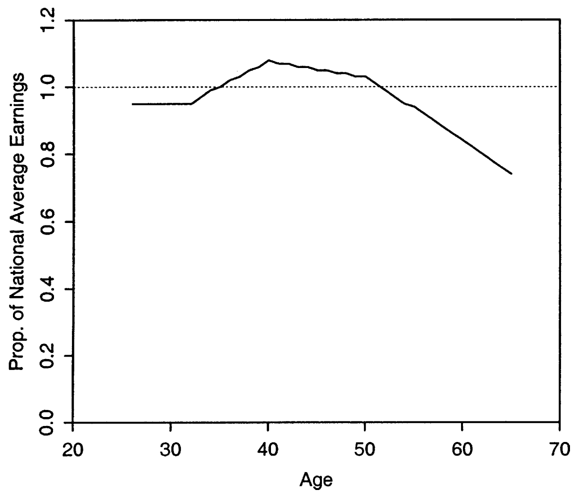

Fig. 1. Life-time earnings profile.

'typical' UK worker is different at different ages due to cohort effects and the relationship between early retirement and salary.

We will assume that the wage profile remains constant over time, but is subject to annual uprating in line with the real growth rate in national average earnings. Real earnings are assumed to have a mean growth rate of 2% p.a. (equal to the average real growth rate over the past 50 years), but subject to some volatility. We adopted a fixed moments model.

3.1.4. Unemployment

Unemployment is modelled as a binary variable (1: employed, 0: unemployed) for each period, with the age-dependent probability of unemployment taken as equal to the 1998–1999 UK national average unemployment rate for men of that age.[20,21] We also allowed for some stochastic variation around these basic rates to reflect the fact that unemployment rates vary over time.

3.2. Control variables

The model incorporates three control variables: variables that must be set by either the pension plan member or the provider for each period in the model.

3.2.1. Contribution rate

The first of these is the pension fund contribution rate. We have assumed a constant contribution rate of 10% of earnings, which is typical of DC plans in the UK. A more general analysis could consider the possibility that the contribution rate might vary over time in response to changes in salary levels and to performance of the plan assets.

[20] To be precise, we took Office for National Statistics data for the unemployment rates of men of various age ranges over the period from December 1988 to February 1999. Interpolations were then taken to approximate the unemployment rates of men of specific ages, and the resulting unemployment rates were taken as proxies for the likely unemployment probabilities of our member as he ages over time.

[21] An alternative measure of not contributing to the pension plan would be related to the activity rate, with the probability of a labour market interruption at any age measured by one minus the activity rate at that age. This gives a lower figure than one minus the unemployment rate, reflecting the fact that many people are not in work for other reasons than being unemployed, e.g., child care, disability, illness, etc.

3.2.2. Asset allocation

The key control variable in the model is the asset allocation: that is, the asset composition of the pension fund portfolio. We consider the following five asset-allocation strategies:

1. a '50/50' low-risk strategy that is 50% in T-bills and 50% in bonds (for most asset-return models this was found to be the minimum-risk static strategy);
2. a 'pension-fund-average' (PFA) strategy equal to the average allocation of UK pension funds in 1998 (namely 5% in T-bills, 51% in UK equities, 15% in UK bonds, 5% in UK property, 20% in international equities, and 4% in international bonds):[22] this might be considered a high-risk strategy on account of its high equity weighting;
3. a 'lifestyle' strategy with a 100% weighting in the PFA portfolio for the first 30 years, followed by a 10% p.a. switch into the 50/50 portfolio during the remaining 10 years of the accumulation plan;
4. a 'threshold' (or 'funded status') strategy (e.g., see Derbyshire, 1999), which is:
   - 100% invested in the PFA portfolio if the current pension ratio[23] is below a lower threshold (T_L),
   - 100% in the 50/50 portfolio if the current pension ratio is above an upper threshold (T_U), and
   - linearly increases in the 50/50 portfolio as the current pension ratio rises from T_L to T_U.

   In this study we concentrate on the thresholds $T_L = 0.4$ and $T_U = 0.8$. We also report briefly on results for different thresholds;
5. a 'portfolio insurance' strategy, based on the constant proportion portfolio insurance (CPPI) framework of Perold and Sharpe (1988), Black and Jones (1988) and Black and Perold (1992) in which the weight in the high-risk portfolio is given by

   $$\text{Weight in the PFA portfolio} = C_M(1 - C_F(\text{Floor/Fund})) = C_M(1 - C_F(\text{Liabilities/Fund})),$$

   where C_F is a parameter measuring the significance attached to the fund being above a floor (where in this example the floor is set at the level of the liabilities in a comparable DB plan[24]), and C_M is the multiplier attached to the quasi-surplus ratio: CPPI strategies have values of C_M exceeding unity. The remaining proportion of the fund is invested in the low-risk 50/50 portfolio.[25] We concentrate in this study on the values $C_F = 0.5$ and $C_M = 2$. We also report briefly on the effect of different values of C_F and C_M.

The first two asset allocations are static. The third strategy is dynamic, but without any feedback. The fourth strategy incorporates a simple form of feedback control. If investment performance has been poor so that the current value of the fund is less than 40% of that needed for the member to retire immediately, the fund is fully invested in the PFA portfolio in order to benefit from the higher expected returns. If, on the other hand, investment performance has been good so that the value of the fund is above 80% of the current pension ratio, the fund should be fully invested in the 50/50 portfolio to protect its value. If the current pension ratio is 50%, the fund should be 25% invested in the 50/50 portfolio (that is, $(0.5 - T_L)/(T_U - T_L)$), etc. The fifth strategy also involves feedback control, but the weight in the PFA portfolio moves in the opposite direction to that of the threshold strategy and rises in line with the quasi-surplus ratio in the fund.[26]

[22] Source: the CAPS performance measurement service.

[23] Here we take the 'current pension ratio' at time t to be an approximate immediate pension $(F(t)/\bar{a})$ divided by $\frac{2}{3}$ of the member's current salary, where $F(t)$ is the current fund size and $\bar{a}$ is an *average* annuity factor for retirees at age 65.

[24] Recall that the ratio of the liabilities to the fund value at retirement is equal to the inverse of the pension ratio.

[25] Short selling was not allowed and the portfolio weights were restricted to lie in the range 0–100%.

[26] These two strategies are consistent with different investor attitudes to risk and return. The threshold strategy gives no credit for pensions exceeding the upper threshold but does not penalise underperformance strongly. In contrast CPPI does give credit for outperformance above the adjusted floor and penalises underperformance severely.

All these strategies are simple both to explain and to implement and hence consistent with the underlying theme of the paper. Optimal dynamic asset-allocation strategies can, of course, be used, but are much less straightforward to explain (see, e.g., Cairns et al., 2000).

3.2.3. Retirement age

Finally, we assume in this version of the model that the retirement age is fixed at 65. A more sophisticated model could allow for the retirement age to depend on, say, the size of the pension fund accumulated, the health of the member, or the likelihood of redundancy triggering early retirement.

4. Simulation output

4.1. Basis of the simulations

The model generates simulated pension ratios corresponding to each of the asset-return models and each of the asset-allocation strategies. Each of these outputs consists of an empirical distribution function plus associated statistics, with the complete set of outputs being generated by 5000 Monte Carlo simulations carried out using specially written VBA routines in Excel.

Value-at-risk statistics have been used to plot the empirical cumulative distribution function of the pension ratio for each model and for each asset-allocation strategy. Selected results are plotted in Figs. 2–8 and we comment in general as follows:

- In each of Figs. 2–8 we have plotted the results for one particular asset-return model.

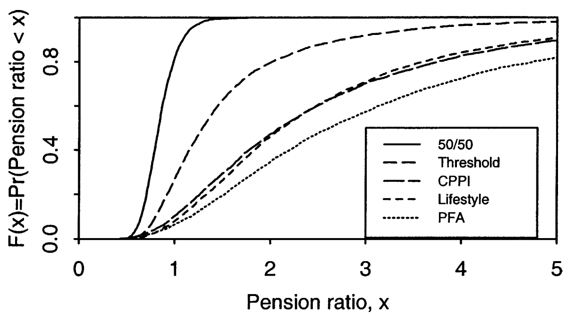

Fig. 2. Multivariate normal model.

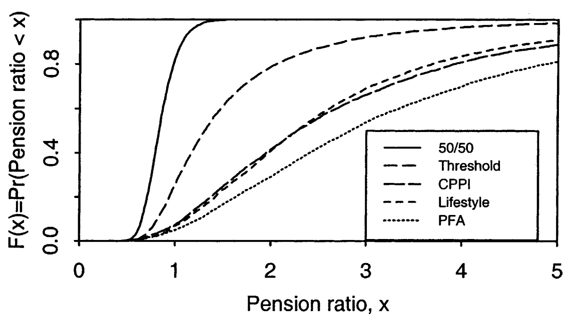

Fig. 3. Mixed multivariate normal model.

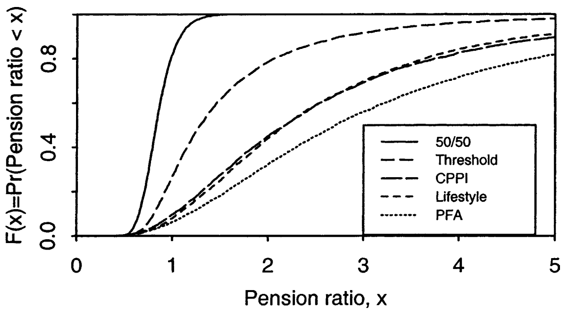

Fig. 4. Multivariate *t* model.

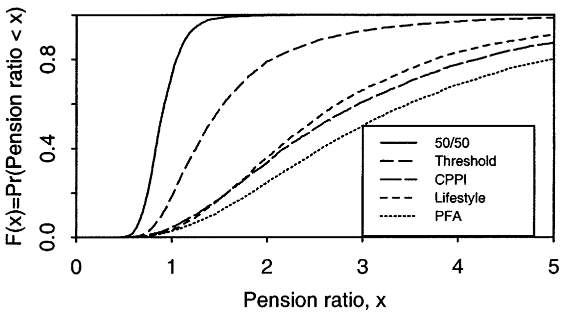

Fig. 5. Multivariate non-central *t* model.

- Each figure contains five curves: one for each asset-allocation strategy considered.
- Each point on a given curve shows the probability that the pension ratio will fall below a particular level. For example, consider the PFA strategy plotted in Fig. 2 (dotted line). We can see that the probability that the pension ratio is less than 2 is approximately 0.35 or 35%. Similarly, the probability that the pension ratio will fall below 1.5 is about 20%. Alternatively we can say that we are 80% confident that the pension ratio will exceed the level 1.5.
- In each figure we can compare the five curves in a qualitative sense:
  - If one curve lies more to the right than the other curves then a particular strategy is likely to deliver higher pensions at retirement than the other strategies.

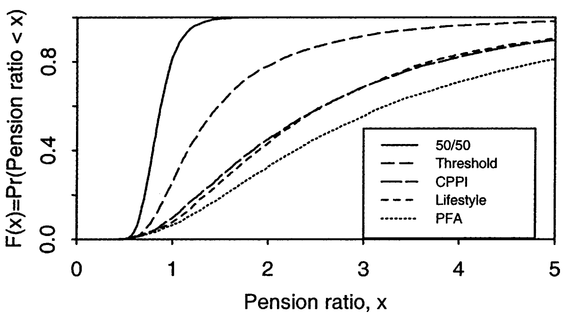

Fig. 6. Bootstrap model.

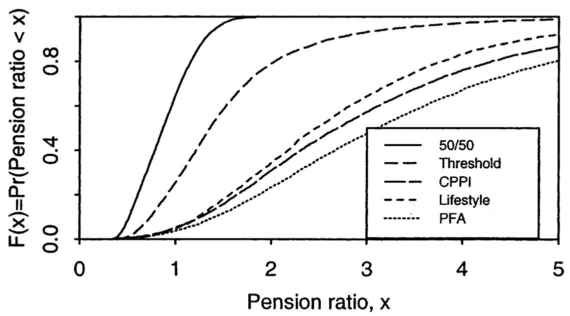

Fig. 7. Markov switching model.

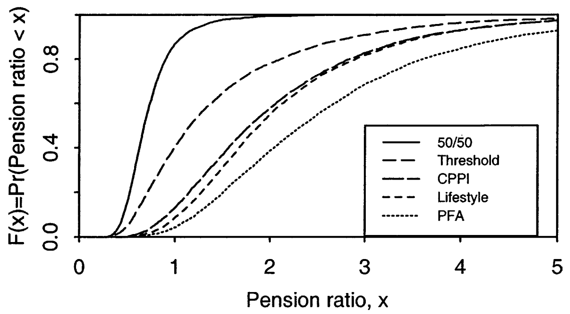

Fig. 8. Wilkie model.

- ○ If a particular curve rises more steeply from 0 to 1 than the others then the strategy is less variable than the others (that is, the pension ratio at retirement is more predictable).

4.2. Further comments on the results

These output graphs are all characteristic of distributions with a single peak. Furthermore, almost all have long tails going off to the right. Perhaps their most striking features are the following:

- They show very considerable dispersion. Consider the very first case (that is, multivariate normal returns and the PFA asset-allocation strategy — Fig. 2, dotted line). The simulated pension ratios vary from about 0.25 at one extreme to 40 at the other (out of 5000 simulations). If one wants an alternative (and more reliable) indicator of their dispersion, the lower 5% quantile is 0.95 and the upper 5% quantile is 8.8. These indicate a 5% chance that the pension ratio will be less than 0.95, and a 5% chance that it will be greater than 8.8.

  Each of the empirical distributions shows a very considerable degree of dispersion however this is measured. These results tell us quite clearly that DC plans can be very risky indeed for members relative to a DB benchmark. [27]

[27] It is legitimate to ask what the equivalent contribution rate would be for a DB scheme. One answer to this is to calculate retrospectively the contribution rate in each simulation run which would deliver the target replacement ratio: that is, the DB rate equals 10% divided by the DC pension ratio. The distribution of these quantities then gives us an indication of the required prospective DB contribution rate: both its mean and dispersion. The mean DB contribution rate turned out to be around 4% for equity based strategies and 12% for the 50/50 portfolio reflecting the latter's low-return/low-risk philosophy. For further discussion, see Cairns (1995).

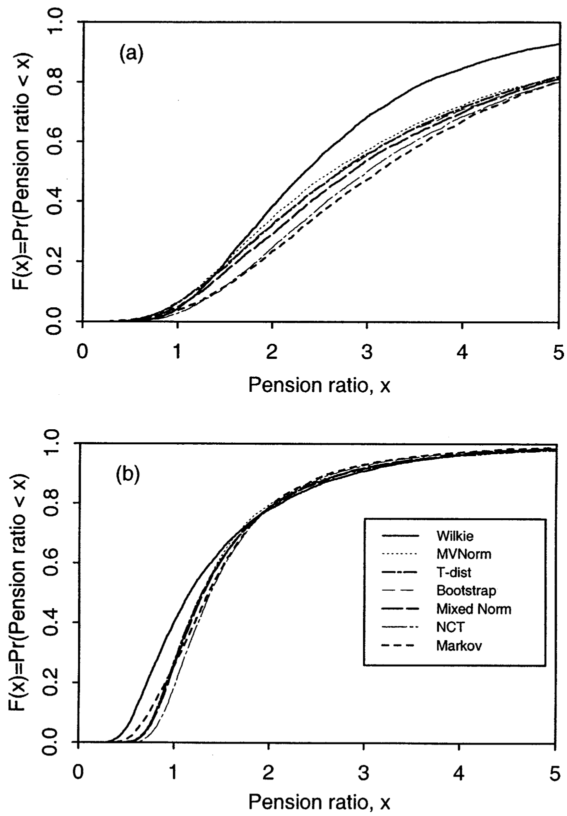

Fig. 9. The effect of model risk on the cumulative distribution function of the pension ratio for two asset-allocation strategies: (a) PFA strategy; (b) threshold strategy ($T_L = 0.4$, $T_U = 0.8$).

- They show very striking asymmetry and kurtosis. For example, in the case just considered, we have a skewness coefficient of 2.4 (that is, a long right-hand tail) and a kurtosis coefficient of 8.9 (fat tails), both of which are very high relative to the normal distribution (0 and 3, respectively). Many of the other combinations have even bigger skewness and kurtosis coefficients. There are also very large differences between the mean and the median, another indication of asymmetry. For the particular case under consideration, which is not untypical, the mean is 3.41 while the median is 2.63.
- Model risk can be assessed by examining the differences between the asset-return models [28] (see Fig. 9). For example, the mean pension ratio across the seven asset-return models for the PFA strategy ranges from 2.70 to 3.73, while for the threshold strategy the range is 1.52–1.66 (see Table 1). Standard deviations for the same two strategies range from 1.53 to 3.21 and from 0.98 to 1.20, respectively. The differences between models become more significant when we consider tail events. For example, the 95% values-at-risk for the same two strategies range from 0.92 to 1.16 and from 0.51 to 0.82, respectively. The range of values for the upper 5% quantiles is rather larger. We can see from Fig. 9 that this range would be rather smaller if we exclude the Wilkie model. This is principally a consequence of the calibration of the Wilkie model to a longer run of data (which had both lower average real returns and lower volatility). The differences we observe are likely to be due mainly to the

[28] See Cairns (2000b) for further discussion of parameter uncertainty and model risk in insurance-related problems and Lee and Wilkie (2000) for an analysis of the differences between different stochastic asset models.

Table 1
Model risk for the PFA and threshold strategies

	Pension ratios for					
	PFA strategy			Threshold strategy		
	Mean	S.D.	95% VaR	Mean	S.D.	95% VaR
Multivariate normal	3.41	2.65	0.95	1.61	1.10	0.74
Mixed multivariate normal	3.56	2.75	1.01	1.63	1.05	0.75
Multivariate *t*	3.49	3.02	0.95	1.64	1.14	0.74
Non-central *t*	3.73	3.21	1.16	1.66	0.98	0.82
Bootstrap	3.51	2.85	0.92	1.64	1.10	0.74
Markov	3.69	2.65	1.11	1.61	1.20	0.64
Wilkie	2.70	1.53	1.04	1.52	1.14	0.51

parameter values used in the Wilkie model rather than due to structural differences between it and the other six models.

- Fig. 9 shows these differences graphically for the same two investment strategies (PFA and threshold). Differences in the parameterisation of the models show up in the way that, e.g., the Markov model always lies a bit to the right of the other curves, while always having roughly the same shape as most of the other curves. The Wilkie model typically stands out of line for reasons discussed above. Apart from the Wilkie model, the results for different models look reasonably consistent with one another. Larger differences show up primarily in the tails which are difficult to see in these graphs. It is in the tails where structural differences between, e.g., the multivariate normal and NCT will be more evident.

  These observations justify our consideration of different asset-return models: that is, we wish to choose asset-allocation strategies that are generally good under most models but, equally well, not especially bad in the remainder of the models. In particular, we should avoid recommending strategies that are very good when we consider one asset-return model, but very bad under another. This is because any model is only an approximation to reality and because the limited amount of historical data we have available still leaves us with a range of plausible models. Nevertheless we regard the evidence as indicating that, except for the extreme tails of the distribution, the asset-return model chosen is not critical. In other words, reassuringly, model risk does not appear to be an important issue.

  A much more important issue is parameter uncertainty. Even if we knew the true model (say, it happens to be the multivariate normal model), we still cannot be sure that our estimates of the parameters of the model are correct because of the limited amount of data. These models have been estimated using data for the last half century: we cannot be at all sure that these estimates will be suitable for conducting simulations over the next 40 years.

  Fig. 9 gives an indication of the possible effects of parameter uncertainty. In particular, consideration of the differences between the Wilkie model and the other six models suggests that parameter uncertainty is rather more significant than model risk.
- Most striking of all, perhaps, is the finding that the simulated pension ratios vary enormously across the different asset-allocation strategies. Moreover, Figs. 2–8 indicate that the effect of the choice of asset-allocation strategy is much more significant than the effect of parameter uncertainty.

  Consider the multivariate normal model (Fig. 2), for example. Table 2 shows that the median pension ratio varies from 0.84 (50/50), through 1.30 (threshold), to 2.63 (PFA). The standard deviation of the pension ratio ranges from 0.18 (50/50), through 1.10 (threshold), to 2.65 (PFA). The asset-allocation strategy is thus obviously extremely important. The huge discrepancy between the two sets of outcomes is also reflected in their percentiles: the 80% VaR, e.g., ranges from 0.71 (50/50), through 0.93 (threshold), to 1.53 (PFA).

Table 2
Asset-allocation strategies with the multivariate normal model

	Pension ratios			
	Median	S.D.	80% VaR	95% VaR
50/50	0.84	0.18	0.71	0.61
PFA	2.63	2.65	1.53	0.95
Lifestyle	2.13	1.86	1.34	0.90
Threshold	1.30	1.10	0.93	0.74
CPPI	2.10	2.06	1.27	0.83

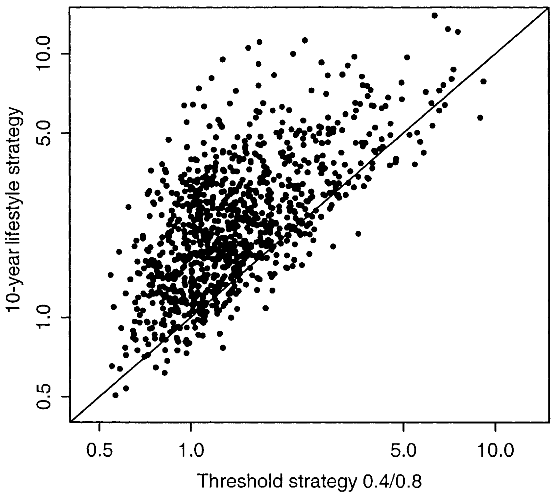

Fig. 10. Multivariate normal model: comparison of 1000 individual scenarios for the threshold and lifestyle models. Points which lie below the diagonal line (142 out of 1000) indicate that the threshold strategy outperformed the lifestyle strategy.

- One tentative conclusion we might draw from Figs. 2–8 is that the most appropriate investment vehicle for a 40-year DC pension plan is, overall, a static, well-diversified, relatively high-return, relatively high-risk portfolio (e.g., the PFA portfolio). None of Figs. 2–8 suggests that a switch into lower-risk assets near retirement is appropriate except possibly as a means of avoiding the most extreme, poor cases. Nevertheless, further investigation of the lifestyle, threshold and CPPI strategies is warranted. [29] For example, the lifestyle strategy can be justified on the grounds that policyholders like the certainty of a reliable estimate of the ultimate pension a few years early even though they may be sacrificing substantial upside potential. Furthermore, all of the asset models considered here have relatively high mean equity returns. More conservative parameter values would change the results to some extent: e.g., by introducing more obvious cross-over points in the left-hand tails.
- In Fig. 2 we can see that the cumulative distribution function for the lifestyle strategy appears to lie to the right of that for the threshold strategy (with $T_L = 0.4$ and $T_U = 0.8$). In Fig. 10 we plot 1000 individual outcomes and see nevertheless approximately 14% of the time the threshold strategy outperforms the lifestyle. It follows therefore that we have something close to (but probably not exact) *first-order* stochastic dominance but not *absolute* stochastic dominance.
- It is useful to investigate how sensitive the threshold and CPPI strategies are to changes in T_L, T_U, C_F and C_M. Sample results are plotted in Figs. 11 and 12. In both cases appropriate choices for the control parameters can

[29] See also Boulier et al. (1999), Deelstra et al. (2000), and Cairns et al. (2000).

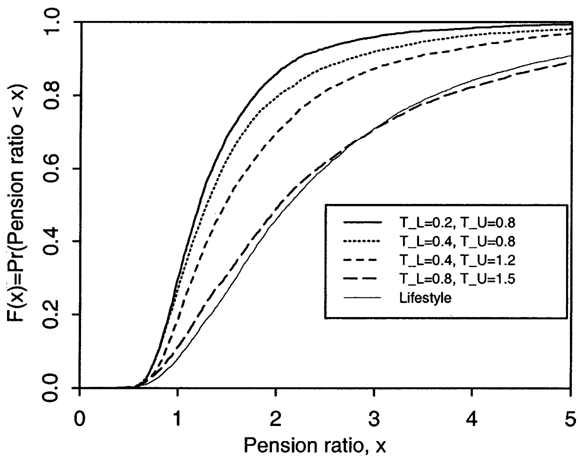

Fig. 11. Sensitivity of the distribution of the pension ratio to changes in the threshold strategy. The lifestyle strategy is plotted for comparison.

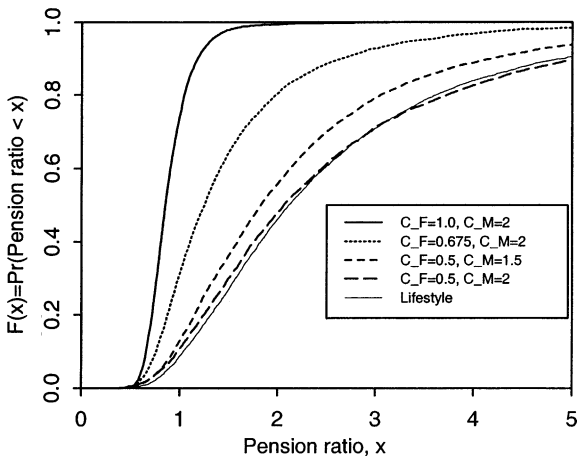

Fig. 12. Sensitivity of the distribution of the pension ratio to changes in the CPPI strategy. The lifestyle strategy is plotted for comparison.

ensure that the fund is always invested 100% in the PFA portfolio or 100% in the 50/50 portfolio. The graphs show that a range of outcomes in between are all possible. This includes results similar to that of the lifestyle strategy. However, as seen in Fig. 10, there will potentially be substantial differences between individual outcomes from following the two strategies.

- A final point to note concerns the accuracy of the statistics quoted in Tables 1 and 2 (that is, how accurate are these estimates of the moments and values-at-risk given we have limited the number of simulations to 5000).

All experiments with 5000 simulations were repeated to check the accuracy of the quoted statistics. For some of the asset-returns models key empirical *moments* of the distribution were found to be poorly determined (that is, did vary significantly from one set of 5000 simulations to the next). The estimate of the mean was well determined in all cases (except, perhaps, for the Markov model [30]). Estimates of the standard deviation were well determined

[30] To illustrate, in two sets of simulations, the Markov model in combination with the PFA strategy had means of 3.69 and 3.75, standard deviations of 2.65 and 3.11, and skewness coefficients of 4.10 and 7.74.

Table 3
VaRs for the threshold strategy

	VaR confidence level			Critical, VaR-CL (%)
	50%	80%	95%	
Multivariate normal	1.30	0.93	0.74	73
Mixed multivariate normal	1.31	0.95	0.75	74
Multivariate *t*	1.30	0.93	0.74	73
Non-central *t*	1.39	1.03	0.82	81
Stationary bootstrap	1.31	0.94	0.74	74
Markov	1.37	0.92	0.64	74
Wilkie	1.17	0.73	0.51	60

except for the mixed multivariate normal and Markov models. In all cases, apart from the NCT and the Wilkie models, the skewness and kurtosis coefficients were poorly determined. The main reason for this sensitivity is clear. The moments (especially skewness and kurtosis) are highly sensitive to the magnitude of the largest one or two outcomes because these are drawn from relatively fat-tailed distributions (the higher the moment, the higher the degree of sensitivity).

In contrast with the estimates of the moments, the value-at-risk statistics were all found to be robust: that is, did not vary significantly from one set of simulations to the next. This is because they do not depend in any way on the extreme outcomes.

These findings support the use of value-at-risk in preference to means and variances in the decision making process, where such decisions are based upon simulation results.

4.3. *Some comparisons between DB and DC schemes*

Recall that the pension ratio is defined as the ratio of the DC pension to a hypothetical benchmark DB pension which is equal to two-thirds of final salary. We now select some particular portfolio allocation strategy, say, the threshold strategy with $T_L = 0.4$ and $T_U = 0.8$, and one or more percentiles or VaR-confidence levels, on which we will make our comparison. Suppose we take these confidence levels to be 50, 80 and 95%.

We then get the results reported in Table 3. If we take the 50% confidence level then our multivariate normal VaR is 1.30, so we would accept that this DC scheme does achieve our target pension ratio of 1. However, if we take the 80% confidence level, our VaR drops to 0.93, and we would reject the DC scheme. At the 95% confidence level the VaR drops further to 0.74. If we want, we can also look at the critical VaR-confidence level — the confidence level at which we would just accept the DC scheme. In this particular case, this critical confidence level is 73%, so we would accept the DC scheme (as achieving our target) if we have a confidence level less than 73%, but reject it if we required a confidence level greater than 73%. The table also lists the comparable results for the other asset-return processes. One can see that these figures do not vary too substantially except at the 95% confidence level and for the Wilkie and NCT models.

We can derive similar tables for other asset-allocation strategies and asset-return models. In most cases the final answer — whether or not the DC plan is an adequate substitute for the DB one — will depend on the precise criterion used (that is, the choice of VaR confidence level) and will also often vary with the asset-return models and the asset-allocation strategy. However for reasonable VaR confidence levels, say in the range 50–80%, the evidence in Table 1 indicates that the asset-return model is much less important than the asset-allocation strategy.

4.4. *The trade-off between risk and contribution rates*

Finally, we can follow the procedure in Section 2 and adjust the contribution rate into the plan until we achieve the desired VaR confidence level. Table 4 shows the contribution rates required by the other asset-allocation strategies

Table 4
Required contribution rates for the multivariate normal model

VaR confidence level for the PFA (%)	99	95	50
VaR for the PFA	0.62	0.95	2.63
	Required contribution rate (%)		
50/50	11.5	15.6	31.3
Lifestyle	9.0	10.6	12.3
Threshold	10.0	12.8	20.2
CPPI	10.3	11.4	12.5

to have the same VaRs as the PFA strategy [31] at the 99, 95 and 50% confidence levels, assuming the multivariate normal model for asset returns. For the lifestyle strategy the contribution rate needed to achieve the 99% VaR for the PFA strategy is actually less than 10%, since the 99% VaR of the PFA is relatively very low (0.62). In all other cases, the required contribution rate to match the PFA is at least as high and often considerably higher. For the 50/50 strategy, a contribution rate of 31.3% is needed to match the PFA median. With the lifestyle strategy, the downside protection achieved by systematically switching into bonds during the last 10 years of the plan requires the contribution rate to rise to 12.3% if the same median pension ratio as the PFA strategy is desired by the plan member. Of the two strategies involving feedback, the threshold strategy looks distinctly more expensive than CPPI except for the most extreme poor outcomes.

5. Conclusions

This paper outlines a simple and very practical methodology for simulating DC pension fund values and key related variables: pension ratios, required contribution rates, and so on. We draw four conclusions from our investigations.

First, we find that DC plans can be extremely risky relative to a DB benchmark (far more so than most pension plan professionals would be likely to admit).

Second, the VaR estimates are very sensitive to the choice of asset-allocation strategy. VaR estimates are also sensitive, but to a lesser extent, to both the asset-returns model used and its parameterisation. The choice of asset-returns model is the least significant of the three, except when considering tail events in asset returns. We tentatively conclude that parameter uncertainty is much less significant than the choice of asset-allocation strategy. However, this is an issue which merits further investigation.

Third, a static asset-allocation strategy with a high equity content delivers substantially better results than any of the dynamic strategies investigated over the long term (40 years) of the sample policy. This is important given that lifestyle strategies are the cornerstone of many DC plans.

Fourth, conservative bond-based asset-allocation strategies require substantially higher contribution rates than more risky equity-based strategies if the same retirement pension is to be achieved.

The model set out here can be fine-tuned and extended in many ways to give practically useful results. For example, we could take account of minimum funding requirements, [32] life-time earnings profiles in different industries, the impact of job histories on individuals' employment and earnings prospects, probabilities of accidents at work, child-rearing leave, personal (biometric) histories (that is, probabilities of disability and ill health), the interrelationship between pensions and other long-term financial contracts (e.g., insurance and life assurance), spouses' pensions, the use of the pension to fund long-term care, and many other issues. Many of these would be no more

[31] Recall that the contribution rate for the PFA is 10%.
[32] For more details on MFRs, see Cairns (1999b, 2000a).

than a fairly simple add-on to our basic model. However, the most important extension is to cover the distribution (post-retirement) phase. This we do in Blake et al. (2000b).

Acknowledgements

We would like to thank Hugh Davies, Nigel Foster, Alexander McNeil, Haris Psaradakis, Ron Smith, Martin Sola, Allan Timmermann and David Wilkie for their advice on various technical aspects of this paper. We would also like to thank Simona Zambelli for research assistance. Sponsorship from the BSI Gamma Foundation is gratefully acknowledged.

Appendix A. Analysis of historical asset returns

Historical returns data were analysed for the period 1947–1998 for six assets: UK T-bills, UK equities, UK bonds, UK property, US equities (denominated in sterling) and US bonds (denominated in sterling). A range of statistics were calculated for this data and their values are presented in Table 5.

Appendix B. Asset-return models

In this appendix we give brief details only of the asset-returns models used in the study since our main purpose is the application of the models rather than an explanation of their development. For further details the reader is referred to Blake et al. (2000a).

In some cases, it might be argued that some of the parameter values look unreasonable (e.g., very high standard deviations for UK equities in state 2 of the mixed multivariate normal and in the Markov switching models). Often in such cases the standard errors of certain parameter estimates are very high (being based upon relatively few observations) with confidence intervals incorporating more reasonable values. These parameter values could be adjusted to more reasonable levels but in a rather arbitrary and subjective fashion. Instead we chose a priori to remain true to the best estimates resulting from the statistical analysis.

Throughout this appendix:

- The asset classes will be numbered as follows: (1) UK T-bills; (2) UK equities; (3) UK bonds; (4) UK property; (5) US equities and (6) US bonds.
- $\zeta_1,\dots,\zeta_6$ will represent *independent* standard normal random variables (i.i.d. $\sim N(0, 1)$).
- $Z_1, \dots, Z_6$ will represent *correlated* standard normal random variables.
- Finally, $r_1, \dots, r_6$ will represent the annual return on the various assets in a given year.
- Where appropriate the random variables above will also be indexed by time, t.

B.1. Stationary moments models

This class of model assumes that the unconditional first and second moments of asset returns are time-invariant and that returns are independent from one year to the next.[33] We have investigated the following models for asset returns.

[33] Furthermore, the first moment should be the arithmetic mean of the series rather than the geometric mean, since as Kolbe et al. (1984) point out: 'The arithmetic mean is the unbiased measure of the expected value of repeated observations of a random variable, not the geometric mean'.

Table 5
Time series properties of total real asset returns and real earnings growth, annual, 1947–1998[a]

	UK T-bills[b]	UK equities[c]	UK bonds[c]	UK property[c]	US equities[d]	US bonds[d]	Earnings[e]
Mean (arithmetic) (%)	1.2788	10.3692	1.5461	4.4783	8.9670	2.1255	2.0908
Standard deviation (%)	4.0455	27.1126	13.9531	10.4522	21.1563	16.9606	2.2074
Standard error (%)	0.5610	3.7598	1.9349	1.4495	2.9339	2.3520	0.3061
Coefficient of variation	3.1635	2.6147	9.0247	2.3340	2.3594	7.9798	1.0558
Correlation matrix[f]							
UK T-bills	1.0						
UK equities	−0.0612	1.0					
UK bonds	0.2563**	0.5441*	1.0				
UK property	0.2720**	0.1854***	0.2016	1.0			
US equities	0.0679	0.4814*	0.2335**	0.0561	1.0		
US bonds	0.2603**	0.1568	0.3046*	−0.0358	0.6818*	1.0	
UK real earnings	0.2110***	−0.0544	−0.3438*	0.3623*	0.0445	−0.0144	1.0
Tests for skewness and kurtosis							
Skewness[g]	−0.8873	1.2085	0.5123	−0.6309	−0.1393	0.9237	−1.0221
Kurtosis[h]	3.6408	8.8360	3.3888	4.3749	3.1399	4.4769	6.2063
Jarque-Bera (normal)[i]	7.7136*	86.4517*	2.6019	7.5458*	0.2105	12.1201*	31.3280*
Expected kurtosis (and d.f.) for *t* [j]	3.67 (13)	9.00 (5)	3.40 (19)	4.50 (8)	3.14 (47)	5.00 (7)	6.00 (6)
Jarque-Bera (*t* with d.f.)[k]	8.3420* (13)	38.1490* (5)	2.5780 (19)	5.2260 (8)	0.1759 (47)	13.3118* (7)	18.2915* (6)
Autocorrelogram at lag[l]							
1	0.786*	−0.210	0.044	0.224	−0.053	−0.013	0.101
2	0.635*	−0.196	0.059	−0.156	−0.226	0.069	−0.074
3	0.564*	−0.124	−0.153	−0.342*	0.087	−0.046	−0.068
4	0.482*	0.198	0.143	−0.251	0.125	−0.046	−0.109
5	0.312*	0.040	0.374	0.073	0.057	−0.192	0.131
6	0.165	−0.235	0.151	−0.009	0.010	0.000	0.021
7	0.050	0.019	0.003	−0.084	−0.072	0.073	−0.164
8	−0.066	0.112	−0.050	0.071	0.044	−0.038	−0.008
9	−0.102	0.124	0.104	0.207	0.205	0.062	−0.090
10	−0.131	−0.071	0.042	0.192	−0.183	0.064	0.058
11	−0.197	−0.173	0.093	−0.033	−0.074	0.097	−0.149
12	−0.210	0.044	−0.061	−0.219	−0.007	−0.130	−0.062
13	−0.197	0.029	0.116	−0.242	0.114	0.019	0.136
14	−0.227	0.075	−0.057	−0.058	0.012	−0.050	0.129
15	−0.258	−0.283*	0.086	0.114	−0.129	−0.056	0.135
16	−0.235	0.099	0.103	0.219	0.028	0.031	−0.001
17	−0.214	0.210	−0.002	0.059	0.128	0.008	−0.036
18	−0.160	0.101	0.061	−0.015	−0.051	−0.056	0.068
19	−0.061	−0.255	−0.149	−0.070	−0.181	0.010	0.167
20	0.004	−0.099	0.036	−0.043	−0.176	−0.129	−0.042
Box-Pierce *Q* statistic[m]	126.03*	33.68*	19.37	34.37*	19.68	7.77	14.05
Cross-correlogram (significant lags)[l]							
UK T-bills	–	–	3, 4	–	–	3	8
UK equities	1	–	–	15	2, 15, 19	7, 12, 19	3
UK bonds	1, 2, 3	–	–	1	–	–	–
UK property	1	1	1, 4	–	–	4, 8	4
US equities	–	6	4	5	–	–	–
US bonds	1	–	–	5, 8	–	–	–

Table 5 (*Continued*)

	UK T-bills[b]	UK equities[c]	UK bonds[c]	UK property[c]	US equities[d]	US bonds[d]	Earnings[e]
Tests for stationarity[n]	1	2	2	1	0	0	1
Difference stationarity (level)[o]	−2.568	−6.745*	−6.295*	−5.472*	−7.156*	−6.873*	−5.206*
Trend stationarity (level)[p]	1.329	0.925	4.083* (0.006355)	0.965	−0.401	0.830	−1.166
Difference stationarity (Δ)[q]	−6.499*	−8.707*	−7.469*	−6.266*	−10.315*	−8.257*	−7.818*
Trend stationarity (Δ)[r]	0.262	0.233	0.275	0.226	0.720	0.646	0.198
Test for mean reversion							
Variance ratio at lag[s]							
2	−1.038	−3.525*	−3.582*	−1.659	−2.905*	−3.790*	−2.757*
3	−1.498	−3.235*	−2.768*	−1.850	−3.464*	−3.173*	−2.789*
4	−1.387	−3.148*	−2.930*	−2.131*	−3.053*	−2.833*	−2.543*
5	−1.115	−2.691*	−2.853*	−2.407*	−2.672*	−2.462*	−2.569*
6	−1.000	−2.323*	−2.467*	−2.211*	−2.412*	−2.400*	−2.311*
7	−0.865	−2.267*	−2.207*	−2.028*	−2.202*	−2.270*	−2.055*
8	−0.744	−2.155*	−2.047*	−2.006*	−2.118*	−2.080*	−2.021*
9	−0.803	−2.033*	−1.998*	−1.980*	−2.058*	−2.003*	−1.891
10	−0.844	−1.877	−1.872	−1.871	−1.837	−1.906	−1.855

[a] Calculated using Eviews (1997).
[b] Real yield (corrected for UK price inflation).
[c] Real total return (yield plus capital gain corrected for UK price inflation).
[d] Real total return (yield plus capital gain, converted into sterling and corrected for UK price inflation).
[e] Real earnings growth rate (UK earnings inflation corrected for UK price inflation).
[f] The test statistic under the null hypothesis of zero correlation is $t=\hat{\rho}_{ij}((T-2)/(1-\hat{\rho}_{ij}{}^2))^{1/2}\sim t(T-2)$ (where $\hat{\rho}_{ij}$ is the sample correlation coefficient between the *i*th and *j*th variables); the 5, 10 and 20% critical values are 2.01, 1.68 and 1.30, respectively, so that correlation coefficients with absolute values in excess of 0.2734, 0.2312 and 0.1808 are statistically significant at these levels, respectively (Mood and Graybill, 1963, p. 358).
[g] The sample skewness coefficient is given by $\hat{\alpha}_3=\hat{\mu}_3/\hat{\mu}_2^{1.5}$, where $\hat{\mu}_j=\sum_t^T(x_{it}-\bar{x}_i)^j/T$.
[h] The sample kurtosis coefficient is given by $\hat{\alpha}_4=\hat{\mu}_4/\hat{\mu}_2^2$.
[i] The Jarque-Bera statistic for the normal distribution under the null hypothesis of no skewness and no excess kurtosis is $\chi_N^2=\frac{1}{6}T\hat{\alpha}_3^2+\frac{1}{24}T(\hat{\alpha}_4-\alpha_4)^2\sim\chi^2(2)$,where the expected kurtosis for the normal distribution is $\alpha_4=3$; the 5% critical value is 5.99 (Jarque and Bera, 1980).
[j] The expected kurtosis for the Student's *t*-distribution is given by $\alpha_4=3+6/(v-4)$, where v (≥5) is the degrees of freedom; what is presented is the kurtosis coefficient (with corresponding degrees of freedom in parentheses) for the *t*-distribution closest to the sample kurtosis coefficient given in footnote h.
[k] The Jarque-Bera statistic for the Student's *t*-distribution with the same degrees of freedom as in footnote j is $\chi_t^2=(1+k)\{\frac{1}{6}T\hat{\alpha}_3^2+\frac{1}{24}T(\hat{\alpha}_4-\alpha_4)^2\}\sim\chi^2(2)$, where $k=[\{\Gamma(2.5)\Gamma\frac{1}{2}(v-4)v^2]/[3\sqrt{\pi}\Gamma\frac{1}{2}v(v/(v-2))^2]\}-1$ (Spanos, 1994).
[l] The test statistic under the null hypothesis of zero autocorrelation or zero cross-correlation $\sim N(0,T^{-1/2})$; the 5% critical value is ±0.28.
[m] The test statistic (where $\hat{\rho}_j$ is the *j*th sample autocorrelation coefficient) under the null hypothesis of no autocorrelation up to *P* lags is $Q=T\sum_{j=1}^P\hat{\rho}_j^2\sim\chi^2(P)$; the critical value for $P=20$ is 31.4 (Box and Pierce, 1970).
[n] Based on the augmented Dickey–Fuller regression $\Delta y_{it}=a_0^i+a_1^it+a_2^iy_{i,t-1}+\sum_{j=1}^Pa_{2+j}^i\Delta y_{i,t-j}+\xi_{it}$; the order of *P* is indicated.
[o] Under the null hypothesis that $y_{it}\equiv x_{it}$ is stationary in differences (and hence generated by a unit root process) against the alternative hypothesis that x_{it} is stationary in levels, which is the one-sided test H_0: $a_2^i=0$, H_1: $a_2^i<0$ the *t*-statistic on c_2^i has a non-standard *t*-distribution given in Dickey and Fuller (1979); the 5% critical value is −3.51.
[p] Under the null hypothesis that $y_{it}\equiv x_{it}$ is stationary about a linear trend (and hence generated by a linear trend) against the alternative hypothesis that x_{it} is stationary in levels, which is the two-sided test H_0: $a_1^i\neq0$, H_1: $a_1^i=0$, the *t*-statistic on c_1^i has a standard *t*-distribution; where H_0 cannot be rejected, the OLS estimate of a_1^i is presented in parentheses.
[q] Tests whether $y_{it}\equiv\Delta x_{it}$ is stationary in differences; the 5% critical value is −3.51.
[r] Tests whether $y_{it}\equiv\Delta x_{it}$ is stationary about a linear time trend.
[s] The variance ratio statistic is defined by

$$VR(i,k)=\left(\frac{\hat{\sigma}_{ik}^2}{\hat{\sigma}_{i1}^2}-1\right)\left[\frac{2(2k-1)(k-1)}{3Tk}\right]^{-1/2},$$

where

$$\hat{\sigma}_{i1}^2=\frac{1}{T-1}\sum_{t=1}^T\left[(x_{it}-x_{i,t-1})-\frac{1}{T}(x_{iT}-x_{i0})\right]^2$$

is the unbiased estimator of ω_i^2 under the assumption that

$$x_{it}=b_0^i+x_{i,t-1}+\xi_{it}$$

and $\xi_{it}\sim N(0,\omega_i^2)$, and where

$$\hat{\sigma}_{ik}^2=\frac{T}{k(T-k)(T-k+1)}\sum_{t=k}^T\left[(x_{it}-x_{i,t-k})-\frac{k}{T}(x_{iT}-x_{i0})\right]^2$$

is an unbiased estimator of the *k*th difference of x_{it} and also an unbiased estimator of ω_i^2 under the assumption that x_{it} follows a random walk. Under the hypothesis of no mean reversion, H_0: $VR(i,k)=0$, the test statistic $VR(i,k)\sim N(0,1)$; the 5% critical value is ±1.96 (Cochrane, 1988; Lo and MacKinlay, 1988, 1989).
* Significant at 5% level.
** Significant at 10% level.
*** Significant at 20% level.

B.1.1. Multivariate normal model

Let $L_1 = [l_{ij}]$ be the lower-triangular Cholesky decomposition of the correlation matrix, C (that is, $l_{ij} = 0$ for all $j > i$, and $C_1 = L_1 L_1'$), μ_i be the mean return on asset i and σ_i be the standard deviation of the return on asset i. Then

$$Z_i = \sum_{j=1}^{6} l_{ij}\zeta_j,$$

$$r_i = \mu_i + \sigma_i Z_i,$$

$$\mu' = (1.28, 10.37, 1.55, 4.48, 8.97, 2.13)\ (\%),$$

$$\sigma' = (4.05, 27.11, 13.95, 10.45, 10.45, 21.16, 16.96)\ (\%),$$

$$L_1 = \begin{pmatrix} 1.0000 & 0 & 0 & 0 & 0 & 0 \\ -0.0612 & 0.9981 & 0 & 0 & 0 & 0 \\ 0.2563 & 0.5608 & 0.7873 & 0 & 0 & 0 \\ 0.2720 & 0.2024 & 0.0233 & 0.9405 & 0 & 0 \\ 0.0679 & 0.4865 & -0.0721 & -0.0629 & 0.8658 & 0 \\ 0.2603 & 0.1731 & 0.1789 & -0.1550 & 0.6735 & 0.6267 \end{pmatrix}.$$

B.1.2. Mixed multivariate normal model

In this model the distribution for the return on asset i is a mixture of two normal distributions with different means and variances. Specifically, for $i = 1, \ldots, 6$, let $I_1, \ldots, I_6$ be independent random variables with $\Pr(I_i = 1) = p_i$ and $\Pr(I_i = 2) = 1 - p_i$ ($I_1, \ldots, I_6$ are also assumed to be independent through time). Then

$$r_i = \begin{cases} \mu_{i1} + \sigma_{i1} Z_i & \text{if } I_i = 1, \\ \mu_{i2} + \sigma_{i2} Z_i & \text{if } I_i = 2, \end{cases}$$

where $Z = L_2\zeta$ and L_2 is the Cholesky decomposition of the correlation matrix described below.

The parameters were estimated as follows. First, the μ_{ik}, σ_{ik} and p_i were estimated using the data summarised in Appendix A by taking each variable in turn and ignoring the other series. Second, the modified correlation matrix $C_2 = L_2 L_2'$ is estimated by matching the model and observed covariances of the r_i.

Since $I_1, \ldots, I_6$ are independent we have

$$E[r_i] = p_i\mu_{i1} + (1 - p_i)\mu_{i2}, \tag{B.1}$$

$$\text{var}[r_i] = (\mu_{i1} - \mu_{i2})^2 p_i(1 - p_i) + p_i\sigma_{i1}^2 + (1 - p_i)\sigma_{i2}^2 \tag{B.2}$$

and for $i \neq j$

$$\text{cov}[r_i, r_j] = (p_i\sigma_{i1} + (1 - p_i)\sigma_{i2})(p_j\sigma_{j1} + (1 - p_j)\sigma_{j2})c_{ij}, \tag{B.3}$$

where

$$c_{ij} = E[Z_i Z_j] = \text{cor}[Z_i, Z_j]. \tag{B.4}$$

Now let $V = [v_{ij}]$ be the target covariance matrix of the r_i. In particular, take $v_{ij} = \rho_{ij}\sqrt{\text{var}[r_i]\text{var}[r_j]}$ where ρ_{ij} is the observed correlation matrix (Table 5). It follows that

$$V = SC_2S + D,$$

Table 6
Parameters of the mixed multivariate normal model

Asset, i	1	2	3	4	5	6
p_i	1.0000	0.9087	0.7018	0.7018	0.8642	0.6268
μ_{i1} (%)	1.28	9.21	5.14	4.31	5.44	4.98
μ_{i2} (%)	–	21.92	−6.92	4.87	31.43	−2.68
σ_{i1} (%)	4.05	18.24	14.57	12.01	20.34	20.26
σ_{i2} (%)	–	66.62	6.21	4.44	4.01	5.45

where

$$S = \text{diag}(p_i\sigma_{i1} + (1 - p_i)\sigma_{i2})$$

and

$$D = \text{diag}([(\mu_{i1} - \mu_{i2})^2 + (\sigma_{i1} - \sigma_{i2})^2]p_i(1 - p_i))$$

$$\Rightarrow C_2 = S^{-1}(V - D)S^{-1}.$$

The fact that C_1 and C_2 are not equal reflects the important observation of Embrechts et al. (1999) that we must always exercise caution when dealing with correlations in a non-normal (strictly non-elliptical) world.

It is straightforward to show that the diagonal elements of C_2 are all equal to 1, so that C_2 is a candidate for a correlation matrix. However, C_2 must also be positive definite. Unfortunately, this is not true in general and, specifically is not true in the present context. This means that there does not exist a correlation matrix C_2 for the Z_i which gives us the desired covariances for the r_i. [34]

As a compromise we must choose a correlation matrix C_2 which gets $SC_2S + D$ as close as possible to V. This is, to some extent, a subjective exercise and we felt that a suitable compromise could be achieved by multiplying the off-diagonal elements of V by the constant factor λ. It was found by trial and error that C_2 was positive definite if λ was less than about 0.89. We chose to take $\lambda = 0.85$ which avoids the potential problems of taking λ right on the threshold of 0.89. [35]

It was found that UK T-bills could be modelled quite adequately without using a mixed-normal model.

The parameters used in the model are given in Table 6 and

$$L_2 = \begin{pmatrix} 1.0000 & 0 & 0 & 0 & 0 & 0 \\ -0.0616 & 0.9981 & 0 & 0 & 0 & 0 \\ 0.2492 & 0.6552 & 0.7132 & 0 & 0 & 0 \\ 0.2454 & 0.2138 & 0.0096 & 0.9455 & 0 & 0 \\ 0.0667 & 0.5658 & -0.1750 & -0.0816 & 0.7988 & 0 \\ 0.2522 & 0.1960 & 0.2053 & -0.1508 & 0.8259 & 0.3886 \end{pmatrix}.$$

[34] This is related to the bivariate log-normal example in Embrechts et al. (1999). They show that the usual correlation coefficient must lie in a range which is narrower than the full range $[-1, +1]$.

[35] In the log-normal example of Embrechts et al. (1999) the maximum correlation coefficient of ρ could be achieved only when the two random variables were perfectly but non-linearly correlated. This implies a reduction in the number of independent variables from two to one. This would be equivalent in the present model to reducing the number of independent variables from 6 to 5 (that is, the value of the sixth variable is defined uniquely by the simulated values of the first five).

B.1.3. Multivariate t model

Here we again use a modified correlation matrix $C_3 = [c_{ij}]$ with Cholesky decomposition L_3. Let $Z = L_3\zeta$. Let $V_1, \ldots, V_6$ be independent chi-squared random variables with $d_1, \ldots, d_6$ degrees of freedom, respectively. Then we take

$$r_i = \mu_i + \sigma_i \frac{Z_i}{\sqrt{V_i/d_i}}\sqrt{\frac{d_i - 2}{d_i}}.$$

Note that this is not the standard definition of a multivariate t-distribution (e.g., see Everitt, 1998). This means that

$$E[r_i] = \mu_i,$$

$$\mathrm{var}[r_i] = \sigma_i$$

and, for $i \neq j$,

$$\rho_{ij} = \mathrm{cor}[r_i, r_j] = \sqrt{(d_i - 2)(d_j - 2)}\, g(d_i) g(d_j) c_{ij},$$

where

$$g(d) = \begin{cases} \sqrt{\pi}\dfrac{(d-2)!}{(d/2-1)!(d/2-1)!2^{d-2}} & \text{if } d \text{ is even,} \\[2ex] \dfrac{1}{\sqrt{\pi}}\dfrac{((d-3)/2)!((d-1)/2)!2^{d-1}}{(d-1)!} & \text{if } d \text{ is odd.} \end{cases}$$

The μ_i and σ_i are the same as in the multivariate normal model. The degrees of freedom $d_1, \ldots, d_6$ were estimated as 13, 5, 19, 8, 47 and 7, respectively.

If we take the target correlation matrix $C_1 = [\rho_{ij}]$ as given, the final relationship in combination with $c_{ii} = 1$ for $i = 1, \ldots, 6$ defines the matrix C_3. Again it does not necessarily follow that C_3 is positive definite (the requirement for it to be a correlation matrix). However, in the present case C_3 is positive definite with Cholesky decomposition:

$$L_3 = \begin{pmatrix} 1.0000 & 0 & 0 & 0 & 0 & 0 \\ -0.0680 & 0.9977 & 0 & 0 & 0 & 0 \\ 0.2661 & 0.6188 & 0.7391 & 0 & 0 & 0 \\ 0.2900 & 0.2300 & -0.0085 & 0.9289 & 0 & 0 \\ 0.0698 & 0.5314 & -0.1477 & -0.0914 & 0.8262 & 0 \\ 0.2798 & 0.1984 & 0.1728 & -0.1771 & 0.7322 & 0.5338 \end{pmatrix}.$$

B.1.4. Multivariate non-central t model

Consider first the univariate model for each asset class, $i = 1, \ldots, 6$. Let V have a chi-squared distribution with d degrees of freedom. Also let δ_i be some real constant. Let

$$Y_i = \frac{Z_i + \delta_i}{\sqrt{V/d}}$$

(in contrast with the multivariate t model, we apply the same V for each i).[36] Then Y_i has an NCT distribution with d degrees of freedom and non-centrality parameter δ_i.

[36] The form used here is the textbook definition of a non-central t-distribution (see Everitt, 1998). The use of independent V_i with different degrees of freedom as with the multivariate t model would have been possible but this would have involved much greater complexity. The form used here has five fewer parameters to estimate.

Now define:[37]

$$\theta=\sqrt{\frac{d}{2}}\frac{\Gamma\frac{1}{2}(d-1)}{\Gamma(\frac{1}{2}d)},$$

then

$$E[Y_i]=\theta\delta_i,$$

$$\mathrm{Var}[Y_i]=\frac{d}{d-2}+\left(\frac{d}{d-2}-\theta^2\right)\delta_i^2$$

and

$$E[(Y_i-E(Y_i))^3]=E[Y_i]\left\{\frac{d(2d-3+\delta_i^2)}{(d-2)(d-3)}-2\,\mathrm{Var}[Y_i]\right\}.$$

Let us now consider the correlations and covariances. Let $C_4=[c_{ij}]$ be the correlation matrix of the Z_i. Then

$$\mathrm{Cov}[Y_i,Y_j]=\frac{d}{d-2}c_{ij}+\left(\frac{d}{d-2}-\theta^2\right)\delta_i\delta_j.$$

Finally, we take $r_i=\exp(m_i+s_iY_i)-1$.

The parameters were estimated as follows. The limited amount of data led to a very wide confidence regions for each parameter: in particular, there was a wide confidence interval for d.[38] The confidence interval for d included $d=4$. This allows us to specify $d=4$ without significantly altering the quality of the fit of the model to the data. The choice of $d=4$ allows us to test the effect of using a fat-tailed distribution.

As with the multivariate t model we estimated parameters for each model separately. The δ_i, given m_i and s_i, were estimated using the method of moments applied to the coefficient of skewness. Finally, the m_i and s_i were estimated using maximum likelihood. This mixed approach is slightly less efficient than full maximum likelihood. However, it was (a) simpler to implement, and (b) left us in a position (see the parameterisation below) to match the between-asset correlations (which was not the case with full maximum likelihood).

In this case

$$(m_1,\dots,m_6)=(0.0270, 0.1272, -0.0048, 0.0726, 0.1205, -0.0166),$$

$$(s_1,\dots,s_6)=(0.0286, 0.1565, 0.1047, 0.0713, 0.1515, 0.1124),$$

$$(\delta_1,\dots,\delta_6)=(-0.3169, -0.2991, 0.0626, -0.3491, -0.2271, 0.1511),$$

$$C_4=L_4L_4',$$

[37] $\Gamma(k)$ is the gamma function under which $\Gamma(k)=\int_0^\infty x^{k-1}\,\mathrm{e}^{-x}\,\mathrm{d}x$. If k is a positive integer then $\Gamma(k)=(k-1)!$.

[38] Instead of maximising over all parameters we can fix d at a value other than the maximum-likelihood estimate and then maximise the log-likelihood over the remaining parameters. If the maximum of this restricted log-likelihood function lies within $\frac{1}{2}\chi_1^2(0.95)=1.92$ of the true maximum log-likelihood then the fixed value d lies within an approximate 95% confidence interval (see, e.g., Silvey, 1975).

where the Cholesky matrix:

$$L_4 = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0.2531 & 0.9674 & 0 & 0 & 0 & 0 \\ 0.3110 & 0.1864 & 0.9320 & 0 & 0 & 0 \\ 0.6221 & 0.1786 & -0.1748 & 0.7420 & 0 & 0 \\ 0.2673 & 0.0639 & -0.0179 & 0.5257 & 0.8049 & 0 \\ 0.2779 & -0.0775 & 0.2064 & 0.1374 & 0.6804 & 0.6264 \end{pmatrix}.$$

(As with the multivariate t model we are fortunate that C_4 is positive definite allowing us to match the target correlation matrix for the r_i.)

B.1.5. Bootstrap model

This model was developed by Efron (1979) and Politis and Romano (1994) and allows the sampling distribution of any stationary weakly dependent time series to be approximated non-parametrically. The procedure involves repeated resampling with replacement from the original time series of asset returns to create pseudo-time series of returns that by construction are drawn from the true but unknown distribution of asset returns, so long as that distribution has stationary moments. Because the time series are weakly dependent, the resampling is performed using blocks of consecutive observations. To achieve full stationarity, the number of observations in the block should be random and follow a geometric distribution; in addition the data is 'wrapped around a circle' so that the first observation in the series follows the last. A variation on this is the moving block technique in which the block size is fixed. Further, if there are a number of series and these are contemporaneously correlated, this correlation structure is preserved by extending the blocks horizontally and using the same dated observations for the second and subsequent series as used in the blocks for the first series.

B.2. Regime-switching model: Markov switching model

Recall again the mixed multivariate normal model. It was assumed that the state variables $I_1, \dots, I_6$ were independent of the values taken in previous years. This is the simplest example of a Markov switching model (e.g., Hamilton, 1994, Sections 22.2 and 22.4).

Instead we can allow the state occupied in each year to be dependent upon the previous state.[39] It was found that two states were sufficient with no dependence between the state-generating processes. Thus we define the transition probabilities $\pi^i_{jk} = \Pr(I_i(t) = k | I_i(t-1) = j)$. Other than this time dependence, the returns model is the same as for the mixed multivariate normal model, but with different parameter values. For notational purposes we write $Z = H_5\zeta$ using the Cholesky decomposition, L_5, of a modified correlation matrix C_5.

Let p_i be the unconditional probability that asset i is in state 1 in a given year. Then $p_i = \pi^i_{12}/(\pi^i_{12} + \pi^i_{21})$. The unconditional means and variances for the $r_i(t)$ are then given by Eqs. (B.1)–(B.4).

[39] For an application of Markov switching models to Australian data see Harris (1999).

Table 7
Parameters of the Markov switching model

Asset, i	1	2	3	4	5	6
π^i_{11}	0.9494	0.9722	0.7590	0.9136	0.9385	0.9408
π^i_{22}	0.9208	0.6733	0.4942	0.8568	0.4195	0.0000
p^i_1	0.6105	0.9215	0.6773	0.6238	0.9043	0.9441
μ_{i1} (%)	3.62	10.28	5.51	4.79	6.69	−0.58
σ_{i1} (%)	1.99	17.92	14.66	12.89	20.82	12.85
μ_{i2} (%)	−2.86	11.33	−6.74	3.99	29.96	47.94
σ_{i2} (%)	3.26	71.61	6.30	3.97	1.72	6.33

The parameter values are given in Table 7 with

$$L_5 = \begin{pmatrix} 1.0000 & 0 & 0 & 0 & 0 & 0 \\ -0.0963 & 0.9953 & 0 & 0 & 0 & 0 \\ 0.3906 & 0.6566 & 0.6542 & 0 & 0 & 0 \\ 0.3940 & 0.2349 & -0.1580 & 0.8744 & 0 & 0 \\ 0.0988 & 0.5228 & -0.2201 & -0.1621 & 0.8014 & 0 \\ 0.4619 & 0.2486 & 0.0589 & -0.3130 & 0.7513 & 0.2428 \end{pmatrix}.$$

B.2.1. Fundamentals model: Wilkie model

A fundamentals model postulates that an asset's return is explained by fundamental economic and asset-related variables such as interest rates, inflation and dividend yield. The Wilkie (1995) model is an example of such a model that is widely used in the UK actuarial profession.

The Wilkie model consists of a system of mean-reverting equations, each with NIID errors and estimated on annual data over the period 1923–1994. Full details of the model and parameter values used in this paper can be found in Wilkie (1995). The model requires different data from the other models analysed here (including dividend yields and interest rates rather than asset returns) so it was not possible to fit the Wilkie model to the same data as the other models.

This means that the Wilkie model gives rather different results from the other models as discussed in the main text.

References

Adams, C., 1999. Older people find employers put youth before experience. The Financial Times, January 17.

Artzner, P., Delbaen, F., Eber, J.-M., Heath, D., 1997. Thinking coherently. Risk 10 (11), 68–71.

Artzner, P., Delbaen, F., Eber, J.-M., Heath, D., 1999. Coherent measures of risk. Mathematical Finance 9, 203–228.

Black, F., Jones, R., 1988. Simplifying portfolio insurance for corporate pension plans. Journal of Portfolio Management 14 (4), 33–37.

Black, F., Perold, A., 1992. Theory of constant proportion portfolio insurance. Journal of Economic Dynamics and Control 16, 403–426.

Blake, D., 1996. Efficiency, risk aversion and portfolio insurance: an analysis of financial asset portfolios held by investors in the United Kingdom. Economic Journal 106, 1175–1192.

Blake, D., Orszag, J.M., 1997. Portability and preservation of pension rights in the United Kingdom. In: Report of the Director General's Inquiry into Pensions, Vol. 3. Office of Fair Trading, London.

Blake, D., Cairns, A., Dowd, K., 2000a. PensionMetrics I: stochastic pension plan design and value-at-risk during the accumulation phase. In: Proceedings of the Third Annual BSI Gamma Foundation Conference on Global Asset Management, Lugano, November 1999. BSI-Gamma Working Paper 19.

Blake, D., Cairns, A., Dowd, K., 2000b. PensionMetrics II: stochastic pension plan design and utility-at-risk during the distribution phase. In: Proceedings of the Fourth Annual BSI Gamma Foundation Conference on Global Asset Management, Rome, October 2000. BSI-Gamma Working Paper 20.

Boulier, J.-F., Huang, S.-J., Taillard, G., 1999. Optimal management under stochastic interest rates: the case of a protected pension fund. In: Proceedings of the Third IME Conference, London, Vol. 2, pp. 1–23.
Box, G.E.P., Pierce, D.A., 1970. Distribution of residual autocorrelations in autoregressive integrated moving average time series models. Journal of the American Statistical Association 65, 1509–1526.
Brown, S., Goetzmann, W., Ross, S., 1995. Survival. Journal of Finance 50, 853–873.
Cairns, A.J.G., 1995. Pension funding in a stochastic environment: the role of objectives in selecting an asset-allocation strategy. In: Proceedings of the Fifth AFIR International Symposium, Brussels, Vol. 1, pp. 845–870.
Cairns, A.J.G., 1999a. A multifactor, equilibrium model for the term structure and inflation. In: Proceedings of the Ninth AFIR Colloquium, Tokyo, Vol. 3, pp. 93–113.
Cairns, A.J.G., 1999b. An analysis of the level of security provided by the minimum funding requirement. British Actuarial Journal 5, 585–610.
Cairns, A.J.G., 2000a. Some notes on the dynamics and optimal control of stochastic pension fund models in continuous time. ASTIN Bulletin 30, 19–55.
Cairns, A.J.G., 2000b. A discussion of parameter and model uncertainty in insurance. Insurance: Mathematics and Economics 27, 313–330.
Cairns, A.J.G., Blake, D., Dowd, K., 2000. Optimal dynamic asset allocation for defined-contribution pension plans. In: Proceedings of the tenth AFIR Colloquium, Tromsø, Norway, June 2000, pp. 131–154.
Cochrane, J.H., 1988. How big is the random walk in GNP? Journal of Political Economy 96, 893–920.
CSFB, 2000. Equity Gilt Study. Credit Suisse First Boston, London.
Deelstra, G., Grasselli, M., Koehl, P.-F., 2000. Optimal investment strategies in a CIR framework. Journal of Applied Probability 37, 936–946.
Derbyshire, G., 1999. Strategic decisions for DC. Pensions Management, October, pp. 56–57.
Dickey, D., Fuller, W.A., 1979. Distribution of the estimates for autoregressive time series with a unit root. Journal of the American Statistical Association 74, 427–431.
Dowd, K., 1998. Beyond Value at Risk: The New Science of Risk Management. Wiley, Chichester.
Efron, B., 1979. Bootstrap methods: another look at the jackknife. Annals of Statistics 7, 1–26.
Embrechts, P., McNeil, A.J., Straumann, D., 1999. Correlation: pitfalls and alternatives. Risk 12 (5), 69–71.
Everitt, B.S., 1998. The Cambridge Dictionary of Statistics. Cambridge University Press, Cambridge.
Eviews, 1997. User's Guide Version 3. Quantitative Micro Software, Irvine, CA.
Flesaker, B., Hughston, L.P., 1996. Positive interest. Risk 9 (1), 46–49.
Goetzmann, W., Jorion, P., 1997. A century of global stock markets. Working Paper 5901. National Bureau of Economic Research, Cambridge, MA.
Hamilton, J.D., 1994. Time Series Analysis. Princeton University Press, Princeton, NJ.
Harris, G., 1999. Monte Carlo Markov chain estimation of regime switching vector autoregressions. ASTIN Bulletin 29, 47–80.
Jarque, C.M., Bera, A.K., 1980. Efficient tests for normality, homoscedasticity, and serial independence of regression residuals. Economics Letters 6, 255–259.
Kolbe, A.L., Read, J.A., Hall, G.R., 1984. The Cost of Capital. MIT Press, Cambridge, MA.
Lee, P.J., Wilkie, A.D., 2000. A comparison of stochastic asset models. In: Proceedings of the tenth International AFIR Colloquium, Tromsø, Norway, June 2000, pp. 407–445.
Lo, A., MacKinlay, A., 1988. Stock market prices do not follow random walks: evidence from a simple specification test. Review of Financial Studies 1, 41–66.
Lo, A., MacKinlay, A., 1989. The size and power of the variance ratio test in finite samples: a Monte Carlo investigation. Journal of Econometrics 40, 203–238.
Merton, R.C., Samuelson, P.A., 1974. Fallacy of the log-normal approximation to optimal portfolio decision making over many periods. Journal of Financial Economics 1, 67–94.
Miles, D., Timmermann, A., 1999. Risk sharing and transition costs in the reform of pension systems in Europe. Economic Policy 29, 253–288.
Mood, A.M., Graybill, F.A., 1963. Introduction to the Theory of Statistics. McGraw-Hill, New York.
Morgan, J.P., 1997. RiskMetrics Technical Document, 3rd Edition. J.P. Morgan, New York.
New Earnings Surveys, 2001. Office for National Statistics. Stationary Office, London.
Perold, A., Sharpe, W., 1988. Dynamic strategies for asset allocation. Financial Analysts Journal, January/February, 16–27.
Politis, D., Romano, J., 1994. The stationary bootstrap. Journal of the American Statistical Association 89, 1303–1313.
Poterba, J., Summers, L., 1988. Mean reversion in stock prices: evidence and implications. Journal of Financial Economics 22, 27–59.
Rutkowski, M., 1997. A note on the Flesaker and Hughston model of the term structure of interest rates. Applied Mathematical Finance 4, 151–163.
Samuelson, P.A., 1963. Risk and uncertainty: a fallacy of large numbers. Scientia 57, 1–6.
Samuelson, P.A., 1989. A case at last for age-phased reduction in equity. In: Proceedings of the National Academy of Science, Vol. 86, Washington, DC, 9048–9051.
Samuelson, P.A., 1991. Long-run risk tolerance when equity returns are mean regressing: pseudoparadoxes and vindication of 'businessman's risk'. In: Brainard, W., Nordhaus, W., Watts, H. (Eds.), Macroeconomics, Finance and Economic Policy: Essays in Honour of James Tobin. MIT Press, Cambridge, MA, 181–200.
Samuelson, P.A., 1992. At last a case for long-horizon risk tolerance and for asset allocation timing. In: Arnott, R.D., Fabozzi, F.J. (Eds.), Active Asset Allocation. McGraw-Hill, Maidenhead.

Siegel, J., 1997. Stocks for the Long Term. Richard D. Irwin, New York.
Silvey, S.D., 1975. Statistical Inference. Chapman & Hall, London.
Spanos, A., 1994. On modelling heteroscedasticity: the Student's *t* and elliptical linear regression models. Econometric Theory 10, 286–315.
Thornton, P.N., Wilson, A.F., 1992. A realistic approach to pension funding. Journal of the Institute of Actuaries 119, 229–312.
Van Eaton, R.D., Conover, J.A., 1998. Misconceptions about optimal equity allocation and investment horizon. Financial Analysts Journal, March/April, 52–59.
Vasicek, O.E., 1977. An equilibrium characterisation of the term structure. Journal of Financial Economics 5, 177–188.
Wilkie, A.D., 1995. More on a stochastic asset model for actuarial use. British Actuarial Journal 1, 777–964.