GLOBAL ASSET LIABILITY MANAGEMENT

M. A. H. DEMPSTER, M. GERMANO, E. A. MEDOVA AND M. VILLAVERDE

[Presented to the Institute of Actuaries, 25 November 2002]

ABSTRACT

Dynamic financial analysis (DFA) is a technique which uses Monte Carlo simulation to investigate the evolution over time of financial models of funds, complex liabilities and entire firms. Although of increasing popularity, the drawback of DFA is the dearth of systematic methods for optimising model parameters for strategic financial planning. This paper introduces *strategic DFA* which employs the only recently mature technology of *dynamic stochastic optimisation* to fill this gap. The new approach is described in terms of an illustrative case study of a joint university/industry project to create a decision support system for strategic asset liability management involving global asset classes and defined contribution pension plans. Although the application of the system described in the paper is to fund design and risk management, the approach and techniques described here are much more broadly applicable to strategic financial planning problems; for example, to insurance and reinsurance firms, to risk capital allocation in financial institutions and trading firms and to corporate investment and business development involving real options. As well as describing the mathematical and statistical models used in the case study, the paper treats econometric estimation, asset return and liability scenario generation, model specification and optimisation, system evaluation and historical backtesting. Throughout the system visualisation plays an important rôle.

KEYWORDS

Dynamic Financial Analysis; Global Capital Markets; Dynamic Stochastic Optimisation; Large Scale Systems; Asset Liability Management; Risk Management; Defined Contribution Pensions; Benchmark Portfolios; Guaranteed Returns

CONTACT ADDRESSES

M A H Dempster, MA, MS, PhD, FIMA, Hon FIA, Centre for Financial Research, Judge Institute of Management, University of Cambridge, Cambridge CB2 1AG, U.K. Tel: +44 1223 339641; Fax: +44 1223 339652; E-mail: mahd2@cam.ac.uk
M Germano, MSc, Pioneer Investment Management Ltd, 5th Floor, 1 George's Quay Plaza, George's Quay, Dublin 2, Ireland. Tel: +353 1 636 4500; Fax: +353 1 636 4600; E-mail: matteo.germano@pioneerinvest.ie
E A Medova, Dip Eng, MA, PhD, Centre for Financial Research, Judge Institute of Management, University of Cambridge, Cambridge CB2 1AG, U.K. Tel: +44 1223 339593; Fax: +44 1223 339652; E-mail: eam28@cam.ac.uk
M Villaverde, MSc, Centre for Financial Research, Judge Institute of Management, University of Cambridge, Cambridge CB2 1AG, U.K. Tel: +44 1223 339651; Fax: +44 1223 339652; E-mail: mv228@cam.ac.uk

> *Dynamic optimization is perceived to be too difficult … It would be nice to have a generic 'sledge hammer' approach for attacking this sort of problem.*
>
> A. D. Smith (1996), p. 1085

1. INTRODUCTION

1.1. *Aims*

1.1.1. Recent years have witnessed the introduction of new investment products aimed at attracting investors who are worried about the volatility of financial markets. The main feature of these products is a minimum return guarantee together with exposure to the upside movements of the markets. While such a return guarantee could be achieved simply by investing in a zero-coupon Treasury bond or similar instrument with expiration equal to the maturity date of the product, this would not allow any expectation of higher returns. Thus there is a need to offer pension products that protect the investor from the downside while maintaining a reasonable expectation of better returns than the guaranteed one.

1.1.2. However, most such current products do not offer a high degree of flexibility; usually, they accept only lump sum investments and have a predetermined maturity of only a few years. This is probably a consequence of the difficulty of reliable long-term forecasting and subsequent determination of the proper asset allocation(s) over the distant time horizon of the investment.

1.1.3. At the same time it is well known that state, and many company, run defined benefit pension plans are becoming inadequate to cover the gap between the contributions of people while working and their pensions once retired. The solution to this problem requires some form of instrument which can fill the gap to allow investors a reasonable income after retirement. A long-term minimum guarantee plan with a variable time-horizon, and in addition to the initial contribution the possibility of making variable contributions during the lifetime of the product, is such an instrument.

1.1.4. Although societally beneficial and potentially highly profitable for the provider the design of such instruments is not a trivial task, as it encompasses the need to do long-term forecasting for investment classes, handling a stochastic number of contributors, contributions and investment horizons, together with providing a guarantee. Stochastic optimisation methodology in the form of dynamic stochastic programming has recently made long strides and is positioned to be the technique of choice to solve these kinds of problems.

1.1.5. This paper describes the approach and outcomes of a joint project between a university financial research centre and a leading firm operating in the European fund management industry to develop a state-of-the-art dynamic asset liability management (ALM) system for pension fund management. The development of this system has been part of an effort undertaken by the firm for the global improvement of its ALM-related technologies and systems.

1.2. *The Pension Fund Problem*

1.2.1. Asset liability management concerns optimal strategic planning for management of financial resources and liabilities in stochastic environments, with market, economic and actuarial risks all playing an important role. The task of a pension fund, in particular, is to guarantee benefit payments to retiring clients by investing part of their current wealth in the financial markets. The responsibility of the pension fund is to hedge the client's risks, while meeting the solvency standards in force, in such a way that all benefit payments are met.

1.2.2. Below we list some of the most important issues a pension fund manager has to face in the determination of the optimal asset allocations over time to the product maturity:-

a) *Stochastic nature of asset returns and liabilities*
   Both the future asset return and the liability streams are unknown. Liabilities, in particular, are determined by actuarial events and have to be matched by the assets. Thus each allocation decision will have to take into account the liabilities level which, in turn, is directly linked to the contribution policy requested by the fund.

b) *Long investment horizons*
   The typical investment horizon is very long (30 years). This means that the fund portfolio will have to be rebalanced many times, making "buy&hold" Markowitz-style portfolio optimisation inefficient. Various dynamic stochastic optimisation techniques are needed to take explicitly into account the on-going rebalancing of the asset-mix.

c) *Risk of under-funding*
   There is a very important requirement to monitor and manage the probability of under-funding for both individual clients and the fund, that is the confidence level with which the pension fund will be able to meet its targets without resort to its parent guarantor.
d) *Management constraints*
   The management of a pension fund is also dictated by a number of solvency requirements which are put in place by the appropriate regulating authorities. These constraints greatly affect the suggested allocation and must always be considered. Moreover, since the fund's portfolio must be actively managed, the markets' bid-ask spreads, taxes and other frictions must also be modelled.

1.2.3. The theory of *dynamic stochastic optimisation* provides the most natural framework for the effective solution of the pension fund ALM problem that will guarantee its users a competitive advantage in the market.

1.2.4. Most firms use static portfolio optimisation models, such as Markowitz mean-variance allocation, which are short-sighted and when rolled forward lead to radical portfolio rebalancing unless severely constrained by the portfolio manager's intuition. Although such models have been extended to take account of liabilities in terms of expected solvency (surplus) levels (see e.g. Mulvey, 1989) these difficulties with static models remain. In practice fund allocations are (thus) wealth dependent and face time-varying investment opportunities, path-dependent returns – due to cash inflows and outflows, transactions costs and time or state dependent volatilities – and conditional mean return parameter uncertainties – due to estimation or calibration errors. Hence *all* conditions necessary for a sequence of myopic static model allocations to be dynamically optimal are violated (see e.g. Scherer, 2002, §1.2).

1.2.5. By contrast, the dynamic stochastic programming models incorporated in the system described below automatically hedge current portfolio allocations against future uncertainties in asset returns and liabilities over a longer horizon, leading to more robust decisions and previews of possible future problems and benefits.

1.3. *Paper Outline*

1.3.1. The next section of the paper sets out the background and basic approach of practical strategic DFA systems for financial planning utilising modern dynamic stochastic optimisation techniques. The remaining sections illustrate these in the context of this case study. Section 3 treats the modelling and econometric estimation of a monthly global asset return model for four major currency areas and the emerging markets which includes macroeconomic variables. In §4, the calibration and stochastic simulation of various versions of this statistical model for use in financial scenario generation for strategic DFA models is discussed. The basic CALM dynamic stochastic optimisation model is treated in §5, including a discussion of risk management objectives, basic constraints, practical constraints and variants of the CALM model for the determination of optimal benchmark portfolios and risk managed return guarantees. Section 6 describes the generation of dynamic stochastic optimisation models for their numerical solution, together with a brief description of solution algorithms and software. Historical out-of-sample backtests of system portfolio recommendations are described in §7 for risk management criteria applied to both terminal fund wealth and the trajectories of the wealth accumulation process. Finally, §8 draws conclusions and indicates directions for future work.

2. STRATEGIC DFA

2.1. *System Design*

2.1.1. Figure 1.1 depicts the processes, models, data and other inputs required to construct a strategic DFA system for dynamic asset liability management with periodic portfolio rebalancing. It should be noted that knowledge of several independent highly technical disciplines is required for strategic DFA in addition to professional domain knowledge. Corresponding to Figure 1.1, Figure 1.2 shows the system design which describes the separate – largely automated and software instantiated – tasks which must be undertaken to obtain recommended strategic decisions once statistical and optimisation models have been specified. Each of the blocks of the latter figure will be treated in detail in a subsequent section of the paper. The outer solid feedback loop recognizes the iterative nature of developing any implementable strategic plan in which process visualisation of data and solutions is key. The inner solid loop will be described in §4. The dotted feedback loops represent possible future developments which will be mentioned in the conclusion.

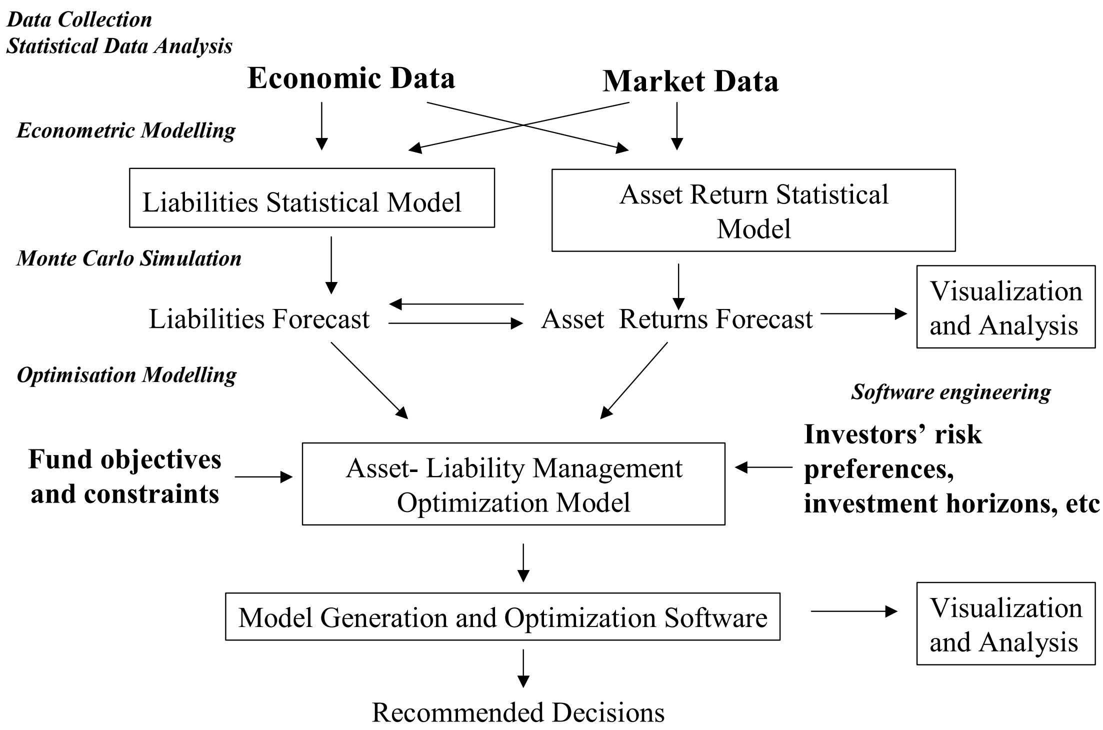

Figure 1.1 Strategic financial planning

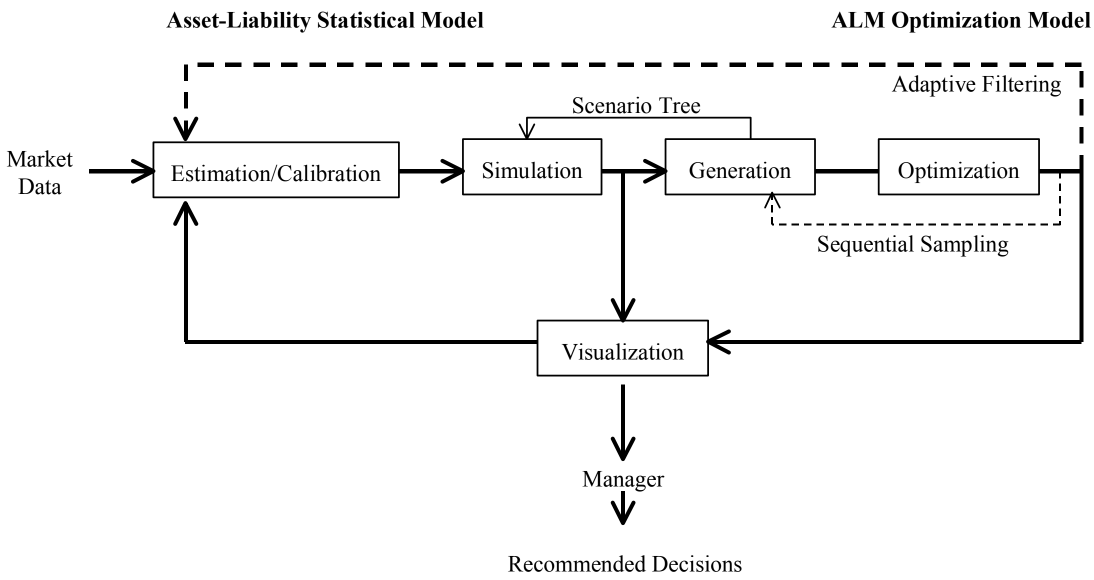

Figure 1.2 System design for strategic financial planning

2.2. *Dynamic Stochastic Optimisation*

2.2.1. As noted above, strategic ALM requires the dynamic formulation of portfolio rebalancing decisions together with appropriate risk management in terms of a dynamic stochastic optimisation problem. Decisions under uncertainty require a complex process of future prediction or projection and the simultaneous consideration of a number of alternatives, some of which must be optimal with respect to a given objective. The problem is that these decisions are only *known* to be optimal or otherwise *after* the realisation of all random factors involved in the decision process. In *dynamic stochastic optimisation* (often termed *dynamic stochastic programming*, as in mathematical programming, see e.g. Dempster (1980)) the unfolding uncertain future is represented by a large number of future scenarios from the DFA simulation process (see e.g. Kaufmann, *et al.* (2001) and the references therein) and contingent decisions are made in stages according to tree representations of future data and decision processes. The initial – *implementable stage* – decisions are made with respect to all possible variations of the future (in so far as it is possible to predict and generate this future) and are thus hedged within the constraints against all undesirable outcomes. This technique also allows detailed 'what-if' analysis of particular extreme future scenarios – forewarned is forearmed!

2.2.2. The methods used are computationally intensive and have only recently become practical for real applications. Each particular optimisation problem is formulated for a specific application combining the goals and the constraints reflecting risk/return relationships. The dynamic nature of stochastic optimisation: decisions – observed output – next decisions – etc … allows a choice of strategy which is the best suited for the stated objectives. For example, for pension funds the objective may be a guaranteed return with a low unexpected risk and decisions reviewed every year. For a trading desk, the objective may be the maximisation of risk adjusted cumulative trading profit with decisions revised every minute, hour or day.

2.2.3. The basic dynamic stochastic optimisation problem treated in this paper is the following. Given a fixed planning horizon and a set of portfolio rebalance dates, find the dynamic investment strategy that maximises the expected utility of the fund's (net) wealth process subject to constraints, such as on borrowing, position limits, portfolio change and risk management tolerances, *viz.*

$$\textit{maximise} \quad \mathbb{E}[U(\mathbf{w}(x))]$$

$$\textit{subject to} \quad A\,x \leqq b.$$

Here U is a specified utility function which is used to express the *attitude to risk* adopted for a particular fund – tailored to broadly match those of its participants over the specified horizon – with regard to the wealth process **w**. (Throughout the paper we use boldface type to represent random entities.) U is used to recommend rebalance decisions which shape the state distributions of **w** over problem scenarios. Risk attitude may concern only *terminal* wealth (Hakansson, 1974; Dempster & Ireland, 1998) or be imposed at each portfolio rebalance date. The (deterministic equivalent form of the) *decision process* x represents portfolio composition at each rebalance date in each scenario subject to the data (A, b) representing the constraints. As such it is a complete contingency plan for the events defined by the scenarios. This basic model will be detailed in §5 and the appendices.

2.3. *Literature Review*

2.3.1. The problem of maximising expected utility under uncertainty subject to constraints can be a highly non-trivial problem. From the point of view of maximising utility the fund will naturally want its set of potential investments to be as large as possible. Thus, it will want the option to invest in global assets ranging from relatively low risk, such as cash, to relatively high risk, such as emerging markets equity. The inclusion of such assets greatly increases the complexity and the amount of uncertainty in the problem since it necessitates the modelling to some degree of not only the asset returns, but also of exchange rates and correlations. Further sources of complexity arise from the multi-period nature of the problem and frictions such as market transaction costs and taxes.

2.3.2. The most well known and probably the most widely used method to solve such a problem is the *mean-variance analysis* pioneered by Markowitz (1952). This analysis can be characterised by a quadratic utility function which depends only on the mean and variance of the portfolio return parameterised by a risk aversion coefficient. Solving the utility maximisation problem for a range of values of the risk aversion parameter gives rise to the *efficient frontier*. This method is now easily implemented in a spreadsheet and only requires an estimate of the mean and covariance of the returns, which are normally obtained from historical data and/or subjective opinion. However, as noted above, the standard implementation of the mean-variance model is static (one-period) and thus fails to capture the multi-period nature of the problem. It also ignores market frictions such as transaction costs. Mean-variance analysis has been extended to incorporate multiple periods and market frictions (see e.g. Steinbach (1999), Horniman *et al*. (2000) and Chellathurai and Draviam (2002)) but at the cost of greatly increased complexity.

2.3.3. In this paper we apply dynamic stochastic optimisation to solve pension fund management problems with global investments. The advance of computing technology and the development of effective algorithms (see e.g. Scott, 2002) have made stochastic optimisation problems significantly more tractable. Following the early work of Bradley & Crane (1972), Lane & Hutchinson (1980), Kusy & Ziemba (1986) and Dempster & Ireland (1988), the growing body of literature concerning the application of stochastic optimisation to fund management problems includes Mulvey and Vladimirou (1992), Dantzig and Infanger (1993), Cariño *et al*. (1994), Consigli and Dempster (1998), Zenios (1998) and Geyer *et al*. (2002) and is a testament to the suitability of this method for solving such problems. A comparison of the application of mean-variance analysis, stochastic control and stochastic optimisation to fund management problems can be found in Hicks-Pedron (1998) where it is shown that dynamic stochastic optimisation performs best in terms of the appropriate Sharpe ratio.

3. ASSET RETURN, EXCHANGE RATE AND ECONOMIC DYNAMICS

3.1. *Asset Return Model*

3.1.1. Our asset return model is in the econometric *estimation* tradition initiated by Wilkie (1986, 1995) and continued, for example, by Cariño *et al.* (1994), Dert (1995), Boender *et al.* (1998) and Duval *et al.* (1999). An alternative approach, in the tradition of Merton (1990), is to set up a continuous time *stochastic differential equation* (sde) model for the financial and economic dynamics of interest, discretise time to obtain the corresponding system of stochastic *difference* equations and *calibrate* the output of their simulation with history by various *ad hoc* or semi-formal methods of parameter adjustment, see, for example, Mulvey & Thorlacius (1998) and Dempster & Thorlacius (1998).

3.1.2. Several other alternative approaches have appeared in the literature which also attempt to generate scenarios known to be arbitrage free within the model. One method widely used for very specific problems in financial stochastic optimisation is sampling scenarios from arbitrage-free lattice paths for the appropriate – e.g. short rate (Zenios, 1998) – arbitrage free model. The resulting sampled scenarios however need not be arbitrage free unless the sampling procedure is carefully controlled (see §4.3). More recently, arbitrage-free methods (Cairns, 2000) and deflator techniques (Smith & Speed, 1998; Jarvis *et al.*, 2001) for designing models in more complex situations have appeared. These modelling approaches involve – at least implicitly – *risk neutral* (i.e. risk discounted) probabilities and *market price of risk premia* to allow simulation of cash flows under real world probabilities. While such approaches are appropriate – indeed necessary – for full discounting for *valuation* purposes, they are totally *inappropriate* for making dynamic *'what-if' forward investment decisions* which must face an approximation of the real world risks. Even for valuation purposes, calibration of complex arbitrage-free models to *current* –but not necessarily past – market data is difficult, not least since the literature on estimating multivariate market prices of risk or state price densities is sparse (but see §3.4 for such a 3-factor yield curve calculation). By contrast with the assumption of no arbitrage – when portfolio decisions are irrelevant to total return (Jarvis *et al.*, 2001) – time varying investment opportunities and potential macro-economic arbitrages occur in the real world.

3.1.3. We have therefore opted for the econometric approach which can – if successful (*cf.* the positive results of system backtests in §7) – model these effects, together with the fact that the estimation procedures involved have been widely employed and most pitfalls in their use documented. Although in our experience some further informal calibration (tuning) of parameter estimates is usually required, for the complex asset return models developed here this has been minimal.

3.1.4. Note that real world scenario generation for stochastic optimisation models by any method may still introduce *spurious* arbitrages due to *sampling* errors. Simple techniques for their suppression will be discussed in §4.3. In this study sampling error has been found to completely swamp statistical parameter estimation error – even assuming that the fitted econometric model actually underlies the data.

3.1.5. Figure 3.1 depicts the global structure of the asset return model involving investments in the three major asset classes – cash, bonds and equities – in the four major currency areas – US, UK, EU and Japan (JP) – together with emerging markets (EM) equities and bonds. Arrows depict possible explanatory dependence.

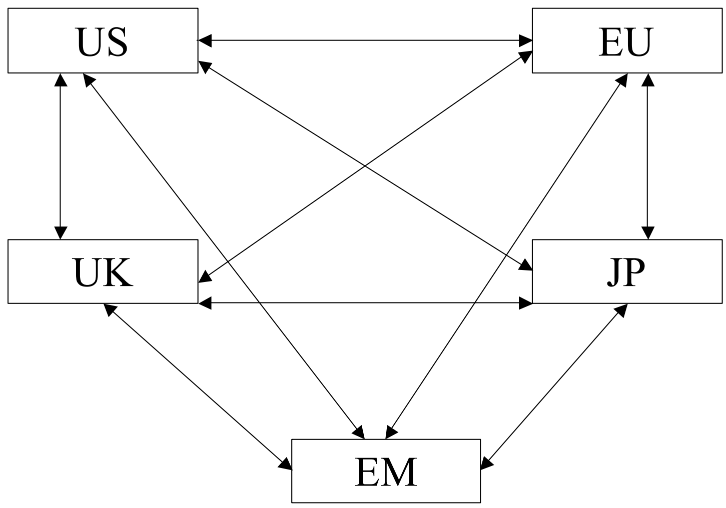

Figure 3.1 Pioneer Asset Return Model Global

3.1.6. Following Dempster & Thorlacius (1998), the approach is to specify a canonical model for each currency area which is linked to the others *directly* via an exchange rate equation and *indirectly* through correlated innovations (disturbance or error terms). For capital market modelling with monthly data this approach was deemed likely to be superior to the usual macroeconomic (quarterly) trade flow linkages (see e.g. Pesaran & Shuermann 2001) between currency areas. Figures 3.2 and 3.3 show respectively at overall and detailed level the structure of the canonical model of a major currency area. Potential liability models in each currency area are shown for completeness although of course pension or guarantee liabilities might be needed only in fewer currencies. The next three sections discuss respectively the canonical model for the capital markets and exchange rate, the emerging markets model and the canonical economic model. The home currency for these models is assumed to be the US dollar, but of course scenarios can be generated in any of the four major currencies since cross rates are forecast and any other currency (e.g. the Euro) can be taken as the home currency for the statistical estimation.

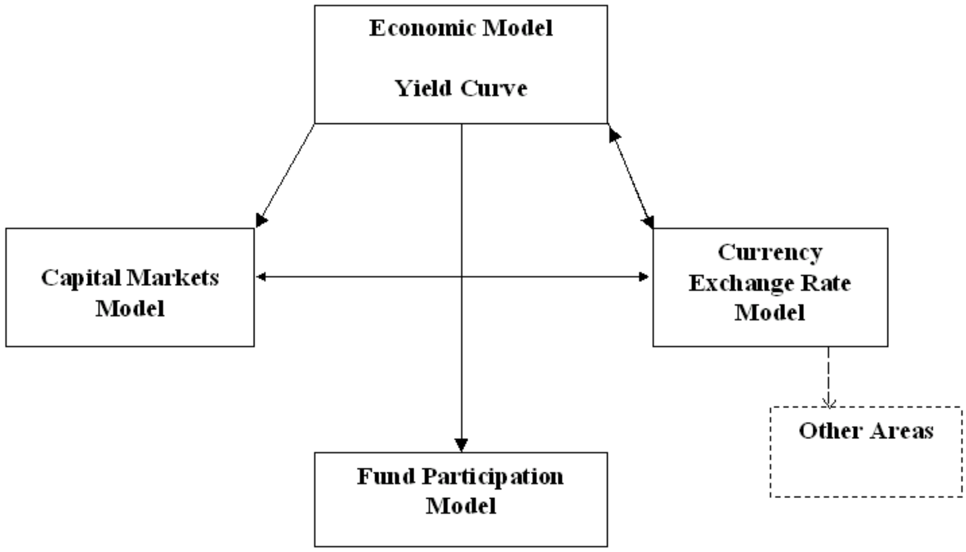

Figure 3.2 Major currency area model structure

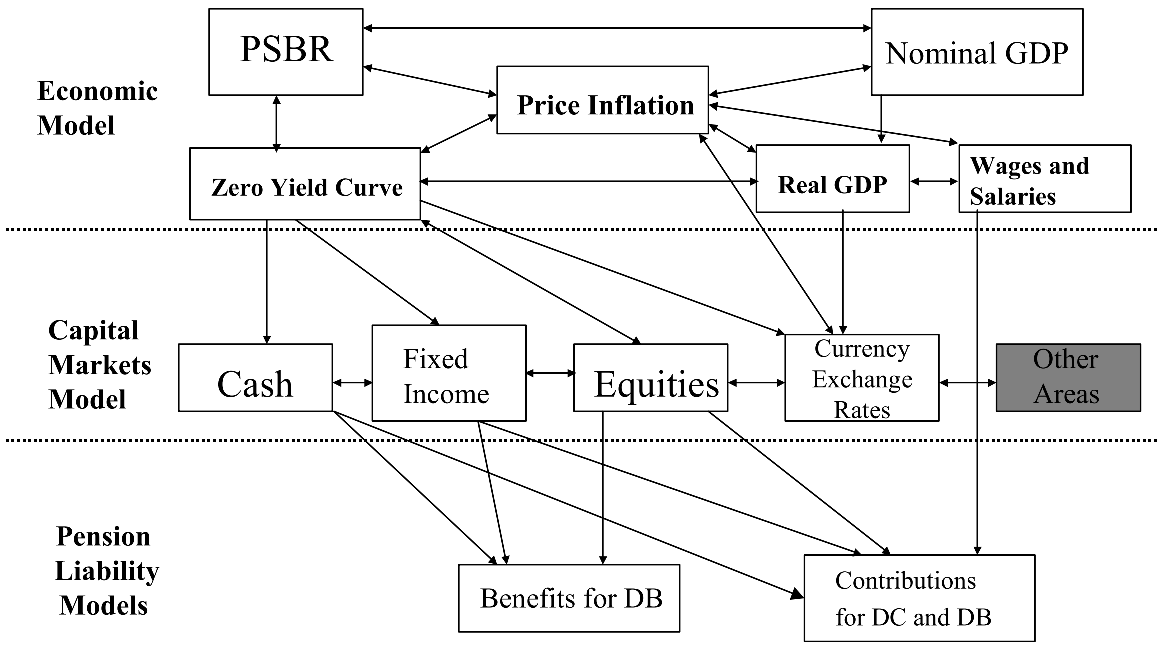

Figure 3.3 Major currency area detailed model structure

3.2. *Capital Markets and Exchange Rate Model*

3.2.1. For simplicity we specify here the evolution of the four state variables – equity (stock market) *index* (**S**), *short* term (money market) interest *rate* (**r**), *long* term (Treasury bond) interest *rate* (**l**) and *exchange rate* (**X**) – in continuous time form as

$$\frac{d\mathbf{S}}{\mathbf{S}} = \boldsymbol{\mu}_s dt + \boldsymbol{\sigma}_s d\mathbf{Z}_s$$
$$d\mathbf{r} = \boldsymbol{\mu}_r dt + \boldsymbol{\sigma}_s d\mathbf{Z}_r$$
$$d\mathbf{l} = \boldsymbol{\mu}_l dt + \boldsymbol{\sigma}_l d\mathbf{Z}_l$$
$$\frac{d\mathbf{X}}{\mathbf{X}} = \boldsymbol{\mu}_X dt + \boldsymbol{\sigma}_X d\mathbf{Z}_X.$$

3.2.2. Here the drifts and volatilities for the four diffusion equations are potentially functions of the four state variables and the d**Z** terms represent (independent) increments of correlated Wiener processes. All dependent variables in this specification are in terms of rates, while the explanatory state variables in the drift and volatility specifications are in original level (**S** and **X**) or rate (**r** and **l**) form.

3.2.3. Detailed specifications of discretised versions of this model are given in Appendix A. The resulting econometric model has been transformed to have all dependent variables in the form of returns and the disturbance structure contemporaneously correlated but serially uncorrelated. In vector terms, the econometric discrete time model has the form

$$\Delta\, \mathbf{x} = \text{diag}(\mathbf{x})\ [\ \mu(\mathbf{x}) + \sqrt{\Sigma}\ \boldsymbol{\varepsilon}\ ],$$

where Δ denotes forward difference, diag (.) is the operator which creates a diagonal matrix from a vector, μ is a first order nonlinear autoregressive filter, $\sqrt{\Sigma}$ is the Cholesky factor of the correlation matrix Σ of the disturbances, and the vector $\boldsymbol{\varepsilon}$ has uncorrelated standardised entries.

3.2.4. Although linear in the drift parameters to be estimated, this model is second order autoregressive and highly nonlinear in the state variables, making its long run dynamics difficult to analyse and potentially unstable. For use in scenario generation over long horizons the model must therefore be linearised so that its stability analysis becomes straightforward. Some linear variants used to date will be discussed in the sequel; we continue to experiment with appropriate forms. Due to its linearity in the parameters this (reduced form) model may be estimated using the *seemingly unrelated regression* (SUR) technique, see e.g. Hamilton (1994) or Cochrane (1997), recursively until a parsimonious estimate is obtained in which all non-zero parameters are statistically significant.

3.3. Emerging Markets Model

3.3.1. After preliminary analysis of the emerging market equity and bond indices (see Table 3.1) using extreme value theory (Kyriacou, 2001) and experimentation with various ARMA/GARCH specifications, it was decided to fit the following ARMA (1,0) model with GARCH (1,1) error structure to index returns *individually, viz.*

$$\mathbf{y}_t = \alpha_0 + \alpha_1 \mathbf{y}_{t-1} - \beta_1 u_{t-1} + \mathbf{u}_t$$
$$H_t = \gamma + pH_{t-1} + qu_{t-1}^2$$
$$\mathbf{u}_t := \sqrt{H_t}\,\boldsymbol{\varepsilon}_t\,,$$

where **y** denotes the (monthly) index return and **ε** is a serially uncorrelated standard normal or student t random variable. Interestingly, although in the EM index data analysed individually the equity index was less extreme than the bond index in terms of tail parameter estimate (Kyriacou, 2001), the above model fit both sets of index data reasonably well with Gaussian innovations. However, these innovations could be expected to be contemporaneously correlated between EM indices and with the innovations of the other variables in the model.

3.3.2. In this case the system model remains as in §3.2, but the enlarged contemporaneous covariance matrix Σ is no longer constant and becomes a process Σ for the entries corresponding to the two extra EM returns. Following a general quasi-likelihood strategy (White, 1982) we may estimate a constant covariance matrix Σ as before using the residuals from the SUR capital market equation estimation and the normalised residuals $u_t / \sqrt{H_t}$ from the individual EM index estimations with sample variance (approximately) 1. Then we compute the Cholesky factor of the corresponding correlation matrix estimate and, for simulation of the full system equation, scale each correlated standardised innovation by the appropriate volatility estimate – constant $\hat{\sigma}$ or time and scenario dependent $\sqrt{\hat{H}_t}$.

3.4. *Economic Model*

3.4.1. In order to capture the interactions of the capital markets with the economy in each major currency area, a small model of the economy was developed with four state variables in nominal values: three financial – consumer price index (CPI), wages and salaries (WS) and public sector borrowing requirement (PSB) – and gross domestic product (GDP). For stability the specification is in terms of returns similar to the capital markets model but with non-state-dependent volatilities, *viz.*

$$\frac{\boldsymbol{CPI}_{t+1} - CPI_t}{CPI_t} = \begin{pmatrix} a_{cpi_1} + a_{cpi_2} CPI_t + a_{cpi_3} WS_t + a_{cpi_4} GDP_t + a_{cpi_5} PSB_t \\ + b_{cpi_2} CPI_{t-1} + b_{cpi_3} WS_{t-1} + b_{cpi_4} GDP_{t-1} + b_{cpi_5} PSB_{t-1} \end{pmatrix} + \sigma_{cpi} \boldsymbol{\varepsilon}_t^{cpi}$$

$$\frac{\boldsymbol{WS}_{t+1} - WS_t}{WS_t} = \begin{pmatrix} a_{ws_1} + a_{ws_2} CPI_t + a_{ws_3} WS_t + a_{ws_4} GDP_t + a_{ws_5} PSB_t \\ + b_{ws_2} CPI_{t-1} + b_{ws_3} WS_{t-1} + b_{ws_4} GDP_{t-1} + b_{ws_5} PSB_{t-1} \end{pmatrix} + \sigma_{ws} \boldsymbol{\varepsilon}_t^{ws}$$

$$\frac{\boldsymbol{PSB}_{t+1} - PSB_t}{PSB_t} = \begin{pmatrix} a_{psb_1} + a_{psb_2} CPI_t + a_{psb_3} WS_t + a_{psb_4} GDP_t + a_{psb_5} PSB_t \\ + b_{psb_2} CPI_{t-1} + b_{psb_3} WS_{t-1} + b_{psb_4} GDP_{t-1} + b_{psb_5} PSB_{t-1} \end{pmatrix} + \sigma_{psb} \boldsymbol{\varepsilon}_t^{psb}$$

$$\frac{\boldsymbol{GDP}_{t+1} - GDP_t}{GDP_t} = \begin{pmatrix} a_{gdp_1} + a_{gdp_2} CPI_t + a_{gdp_3} WS_t + a_{gdp_4} GDP_t + a_{gdp_5} PSB_t \\ + b_{gdp_2} CPI_{t-1} + b_{gdp_3} WS_{t-1} + b_{gdp_4} GDP_{t-1} + b_{gdp_5} PSB_{t-1} \end{pmatrix} + \sigma_{gdp} \boldsymbol{\varepsilon}_t^{gdp} .$$

This is again a second order autoregressive model in the state variables which as shown is *linear* in parameters and *nonlinear* in variables. It may be estimated using the techniques mentioned in §3.2.

3.4.2. With a view to eventually including Treasury bond asset classes of different maturities in the system, a standard 3-factor yield curve model (Campbell, 2000) was developed for fitting to spot yield curve data. The three factors in this model are a very short (one month) rate ($\boldsymbol{R}^0$) and long rate ($\boldsymbol{L}$) corresponding to the capital markets model and a *slope* factor ($\boldsymbol{Y}$: = $\boldsymbol{L}$-$\boldsymbol{R}$) between the short and long rates. By using a time series of monthly yield curve data, it is possible to estimate the evolution over the sample period of the *market prices of risk* (MPRs) for the three factors in volatility units by assuming the model fits the yield curve exactly (commonly referred to as *backing-out* the MPRs).

3.4.3. In more detail, suppose the processes for the three factors $\boldsymbol{R}_0$, $\boldsymbol{Y}$ and $\boldsymbol{L}$ under the real world probabilities satisfy

$$\begin{aligned} d\boldsymbol{R^0} &= (k(\boldsymbol{L} + \boldsymbol{Y} - \boldsymbol{R^0}) + \alpha_{\boldsymbol{R^0}} \sigma_{\boldsymbol{R^0}})dt + \sigma_{\boldsymbol{R^0}} d\boldsymbol{W}_{\boldsymbol{R^0}} := \delta_{\boldsymbol{R^0}} dt + \sigma_{\boldsymbol{R^0}} d\boldsymbol{W}_{\boldsymbol{R^0}} \\ d\boldsymbol{Y} &= (\mu_y - \lambda_y \boldsymbol{Y} + \alpha_y \sigma_y)dt + \sigma_y d\mathbf{W}_y := \delta_y dt + \sigma_y d\mathbf{W}_y \\ d\boldsymbol{L} &= (\mu_L - \lambda_L \boldsymbol{L} + \alpha_L \sigma_L)dt + \sigma_L d\mathbf{W}_L := \delta_L dt + \sigma_L d\mathbf{W}_L . \end{aligned}$$

To calculate market prices of risk time series $\alpha_{R^0 t}$, α_{Yt} , α_{Lt}, for t=1,…,T, we first calibrate the model to detailed yield curve data at t=1 in the usual manner giving estimates of the model parameters $\mu_Y, \mu_L, \lambda_Y, \lambda_L, k, \sigma_{R^0}, \sigma_Y$ and σ_L. Estimates of the real world drifts

$\delta_{R^0_t}, \delta_{Yt}, \delta_{Lt}, t = 1,...,T$ can then be obtained from the historical (monthly) time series for the factors using a suitable backward moving average and data prior to $t = 1$. Estimates of the market prices of risk can then be calculated using the expressions

$$\alpha_{R^0_t} = (\delta_{R^0_t} - kL_t - kY_t + kR^0_t)/\sigma_{R^0}$$
$$\alpha_{Yt} = (\delta_{Yt} - \mu_y - \lambda_y Y_t)/\sigma_Y$$
$$\alpha_{Lt} = (\delta_{Lt} - \mu_L - \lambda_L L_t)/\sigma_L.$$

3.4.4. Figure 3.4 depicts the result of this procedure for the US over nearly a 24 year horizon. Note that while the MPRs of the very short rate and the yield curve slope are highly positively correlated, they are both negatively correlated with the MPR of the long rate, as might be expected for a market which shifts its interest rate risk focus back and forth from short to long term.

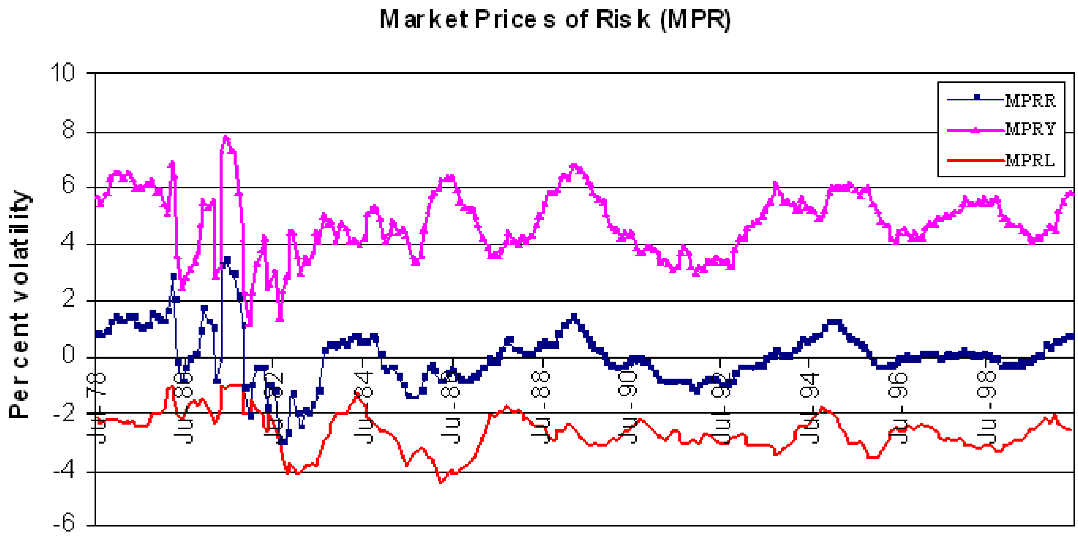

Figure 3.4 Evolution of US yield curve factor market prices of risk

Variable	Corresponding Proxy
S^{US}	S&P 500 stock index
R^{US}	US 3 month T-bill rate
L^{US}	US 30 year T-yield with semi-annual compounding
S^{UK}	FTSE stock index
R^{UK}	UK 3 month T-bill rate
L^{UK}	UK 20 year GILT rate with semi-annual compounding
S^{EU}	MSCI Europe stock index
R^{EU}	German 3 month FIBOR rate
L^{EU}	German 10 year bond yield with annual compounding
S^{JP}	TOPIX stock index
R^{JP}	JP 3 month CD rate
L^{JP}	JP 10 year bond yield with annual compounding
S^{EM}	MSEMEI stock index
B^{EM}	EMBI+ bond index with 10 year average maturity
X^{UK}	UK/US Exchange Rate
X^{EU}	EU/US Exchange Rate
X^{JP}	JP/US Exchange Rate
C^{US}	US CPI
W^{US}	US wage index
G^{US}	US GDP
P^{US}	US public sector borrowing

Table 3.1 Data proxies for model variables

3.4.5. As a preliminary analysis of the interactions of the US macroeconomic and capital market variables over the sample period, these MPRs were regressed on the macroeconomic variables expressed in both levels and returns and significant relationships noted. These accorded well with significant coefficients in the subsequent US system model estimation (see §3.6).

3.5. *Data and System Model Estimation*

3.5.1. Table 3.1 sets out the data used as proxies for the variables of the full system so far discussed. Sources were Data Stream and Bloomberg at monthly frequency from 1977 except for economic variables available only quarterly. Monthly levels were computed for the latter by taking the cube root of the actual quarterly return and finding the corresponding monthly levels between announcements. Figure 3.5 shows equity index evolution in the US and Japan over the 284 month period from July 1977 to February 2001. Dummy variable techniques were required to estimate the effects on constant terms of the bubble and crash period, thereby enabling a meaningful estimation of the Japanese currency area capital market equations. So far they have not proved necessary for recent US history! A consistent database of model data is currently being maintained and updated monthly by the fund manager.

Topix Levels (JP)

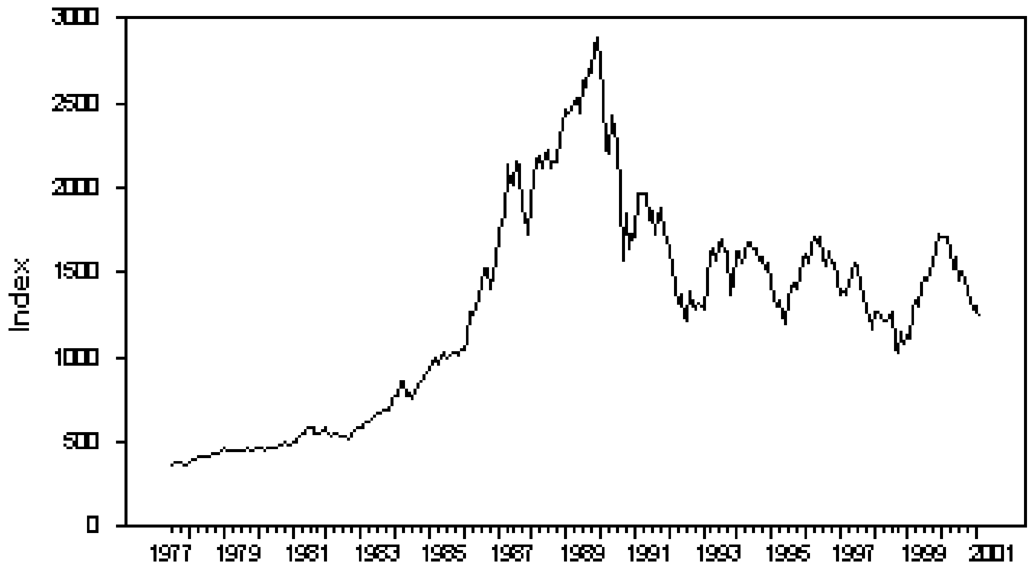

Source: DataStream

Figure 3.5a Equity index evolution in JP

S&P 500 Levels (US)

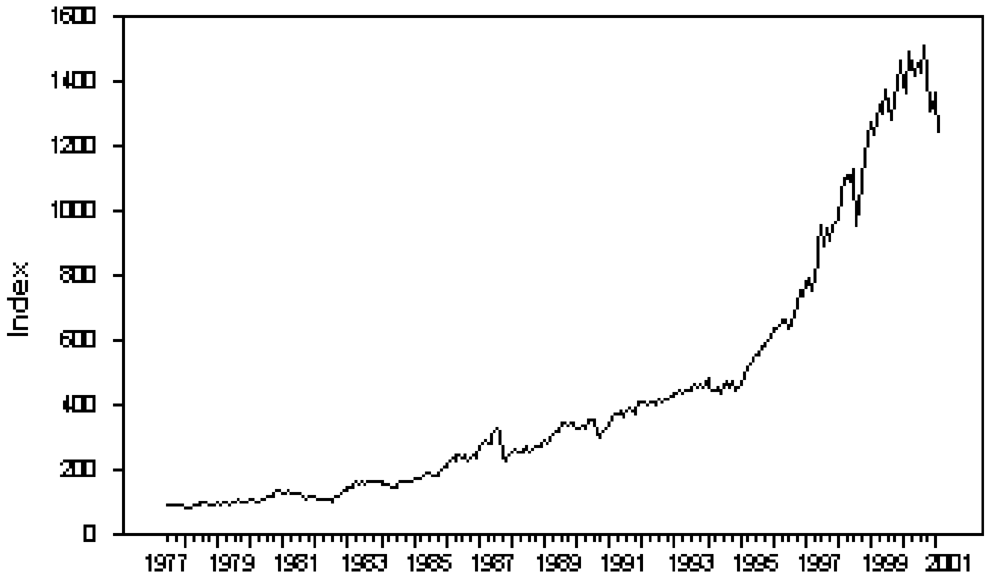

Source: DataStream

Figure 3.5b Equity index evolution in the US

3.5.2. Various subsystems of the full capital markets and economic model have been estimated (see §7) using the SURE model maximum likelihood estimation procedures of RATS (Doan, 1996). For each model the full set of model parameters was first estimated and insignificant (at the 5% level) variables sequentially removed to obtain a parsimonious final model with all statistically significant coefficients. This procedure has been automated in a PERL/RATS script, and (although we are well aware that for given data best variable selection is an NP-hard problem) the automated results agree virtually completely with the much more time consuming hand procedures. Estimation of the emerging market individual ARMA/GARCH equations to yield the AR(1)/GARCH (1,1) specification of §3.4 has been accomplished using S^+. The quasi-likelihood procedure for estimating full models with EM returns was described in §3.4.

3.6. *Results*

3.6.1. We summarise here only illustrative or highly significant findings; more detailed results are forthcoming in Arbeleche (2002).

3.6.2. In this project we have devised a way of presenting econometric model estimation results concisely and graphically. For example, Figure 3.6 shows such an *influence diagram* for a full system model including the US economic variables. Boxes (economic variables) or circles (capital market variables) denote dependent variables (in return form with corresponding adjusted R^2 values shown in percentage terms) and arrows denote a significant influence (solid) or lagged influence (dotted) from a corresponding explanatory variable (tail) to a dependent variable return (tip). The seemingly unrelated regression nature of the model is obvious as each currency area is directly related only through exchange rates and indirectly related through shocks. In light of Meese & Rogoff's (1983a, b) classical view on the inefficacy of macroeconomic explanations of exchange rates even at monthly frequency, after considerable single equation and subsystem analysis we have found that interest rate parity expressed as inter-area short rate differences – *together* with other local capital market variables – has significant explanatory power, while purchasing power parity expressed various ways does not (*cf.* Hodrick & Vassalou, 2002).

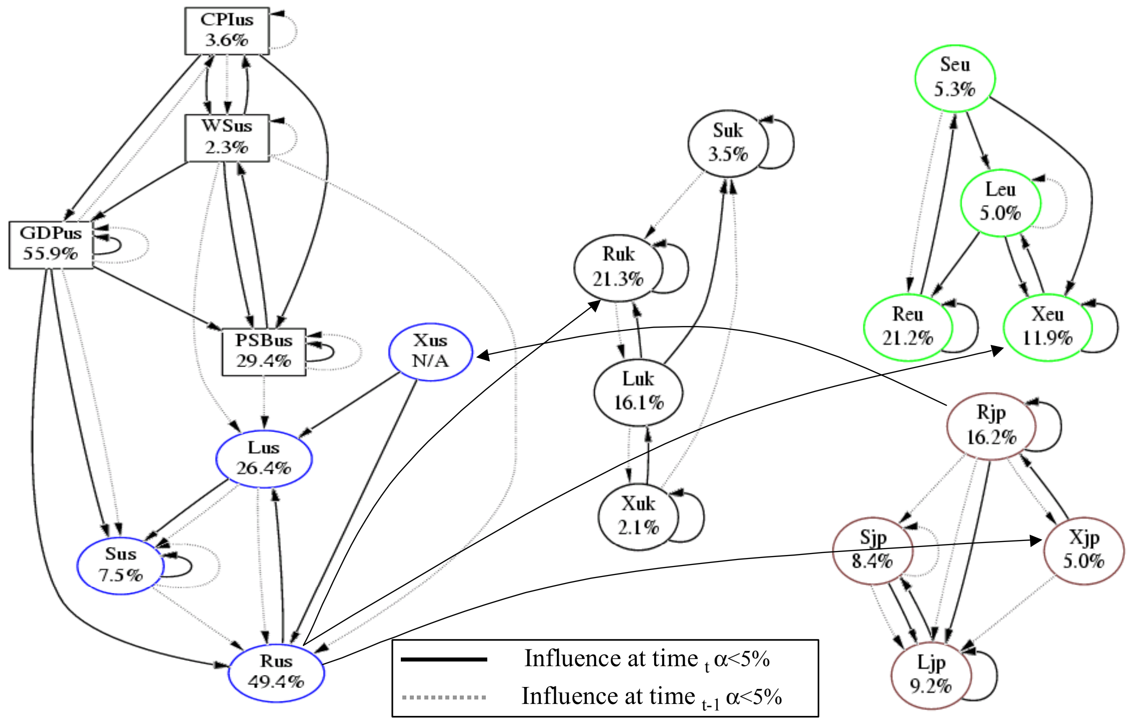

Figure 3.6 Influence diagram for CM +USE + EM 93/01-02/01 system model

3.6.3. Figure 3.7 emphasises our main econometric finding that the world's equity and emerging bond markets and currency exchange rates are linked simultaneously through shocks. The first covariance (diagonal and below)/correlation (above diagonal) matrix is that of raw returns. The second is estimated using residuals from the fitted system model. The circled entries have high correlations and do not change significantly – some actually increase – from the one to the other showing that the dependent variables react mainly to current shocks (innovations) in spite of the stochastic nature of the explanatory variables.

Covariance/correlation matrix of raw returns 93:12 to 01:02

	SUS	RUS	LUS	SUK	RUK	LUK	XUK	SEU	REU	LEU	XEU	SJP	RJP	LJP	XJP	SIE	BIE	CPIUS	WSUS	GDPUS	PSBUS
SUS	**0.002**	-0.084	-0.177	0.719	-0.008	-0.331	-0.023	0.702	-0.078	-0.274	0.202	0.371	0.023	0.075	-0.109	**0.619**	**0.617**	0.217	0.101	0.307	0.175
RUS	0.000	**0.002**	0.420	-0.111	0.224	0.272	-0.182	0.030	0.177	0.345	0.160	0.020	0.096	0.138	0.218	-0.156	-0.152	0.027	0.165	0.207	0.225
LUS	0.000	0.001	**0.001**	-0.086	-0.104	0.599	-0.192	-0.007	-0.001	0.600	0.168	0.232	-0.065	0.080	0.202	0.137	-0.160	-0.076	0.037	0.031	-0.022
SUK	0.001	0.000	0.000	**0.001**	-0.090	-0.373	-0.268	0.762	-0.007	-0.310	0.270	0.347	0.141	0.122	-0.073	**0.583**	**0.566**	0.139	0.006	0.208	0.075
RUK	0.000	0.000	0.000	0.000	**0.001**	0.064	0.277	0.131	0.186	0.021	-0.011	-0.004	0.062	-0.219	0.182	-0.097	-0.001	0.103	0.099	0.072	0.037
LUK	0.000	0.000	0.001	0.000	0.000	**0.001**	-0.020	-0.234	-0.062	0.654	-0.034	0.063	-0.093	-0.053	0.163	-0.035	-0.232	-0.078	0.013	-0.108	-0.045
XUK	0.000	0.000	0.000	0.000	0.000	0.000	**0.000**	-0.225	-0.045	-0.017	-0.585	-0.058	-0.071	-0.051	-0.256	-0.142	-0.188	0.012	0.015	-0.003	0.078
SEU	0.002	0.000	0.000	0.002	0.000	0.000	0.000	**0.003**	-0.034	-0.220	0.437	0.414	0.115	0.174	0.131	**0.554**	**0.529**	0.204	0.103	0.323	0.194
REU	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	**0.002**	0.154	-0.038	-0.118	0.443	-0.083	-0.072	-0.152	-0.070	-0.072	0.041	0.043	-0.069
LEU	0.000	0.000	0.000	0.000	0.000	0.001	0.000	0.000	0.000	**0.001**	-0.178	-0.018	-0.033	0.107	-0.099	-0.020	-0.225	-0.128	-0.009	-0.022	-0.035
XEU	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.001	0.000	0.000	**0.001**	0.271	0.063	0.097	0.369	0.169	0.256	0.082	0.091	0.138	0.067
SJP	0.001	0.000	0.000	0.001	0.000	0.000	0.000	0.001	0.000	0.000	0.000	**0.003**	-0.095	0.180	0.051	**0.490**	**0.391**	-0.004	-0.191	0.069	0.043
RJP	0.000	0.001	-0.001	0.002	0.001	-0.001	0.000	0.002	0.006	0.000	0.000	-0.001	**0.085**	0.088	-0.102	-0.108	0.037	0.097	0.025	0.068	-0.034
LJP	0.000	0.001	0.000	0.001	-0.001	0.000	0.000	0.001	0.000	0.000	0.000	0.001	0.003	**0.015**	-0.172	0.010	-0.120	-0.034	-0.115	0.024	-0.066
XJP	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	-0.001	-0.001	**0.001**	-0.050	0.141	-0.032	0.119	-0.003	0.008
SIE	0.002	0.000	0.000	0.001	0.000	0.000	0.000	0.002	0.000	0.000	0.000	0.002	-0.002	0.000	0.000	**0.005**	0.721	0.051	-0.086	0.038	-0.023
BIE	0.002	0.000	0.000	0.001	0.000	0.000	0.000	0.002	0.000	0.000	0.000	0.001	0.001	-0.001	0.000	0.003	**0.003**	0.243	0.062	0.182	0.155
CPIUS	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	**0.000**	0.519	0.799	0.487
WSUS	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	**0.000**	0.560	0.491
GDPUS	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	**0.000**	0.682
PSBUS	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	**0.001**

Covariance/correlation matrix of residuals 93:12 to 01:02

	SUS	RUS	LUS	SUK	RUK	LUK	XUK	SEU	REU	LEU	XEU	SJP	RJP	LJP	XJP	SIE	BIE	CPIUS	WSUS	GDPUS	PSBUS
SUS	**0.002**	-0.138	-0.148	0.675	0.009	-0.261	0.060	0.674	0.113	-0.241	0.16[illegible]	0.421	-0.017	0.197	-0.134	**0.639**	**0.562**	0.045	0.038	0.129	-0.072
RUS	0.000	**0.001**	0.201	-0.071	-0.004	-0.014	-0.294	-0.003	0.259	-0.007	0.251	-0.027	0.278	0.094	0.319	-0.243	-0.130	-0.347	-0.090	-0.017	0.098
LUS	0.000	0.000	**0.001**	-0.026	-0.169	0.507	0.254	-0.009	0.010	0.449	0.206	0.144	-0.132	-0.058	0.247	0.048	-0.104	-0.062	0.171	0.141	-0.028
SUK	0.001	0.000	0.000	**0.001**	-0.098	-0.340	-0.236	0.732	0.002	-0.325	0.248	0.278	0.152	0.114	-0.104	**0.589**	**0.550**	0.073	-0.073	0.090	-0.105
RUK	0.000	0.000	0.000	0.000	**0.001**	0.103	0.292	0.171	0.138	0.018	0.067	0.175	0.031	-0.229	0.187	-0.098	-0.079	-0.034	0.100	0.031	-0.088
LUK	0.000	0.000	0.000	0.000	0.000	**0.001**	-0.107	-0.120	0.029	0.608	0.106	0.103	-0.170	-0.062	0.277	-0.079	-0.167	-0.140	0.233	0.110	0.061
XUK	0.000	0.000	0.000	0.000	0.000	0.000	**0.000**	-0.184	-0.001	-0.071	-0.609	-0.085	0.067	-0.061	-0.253	-0.143	-0.229	0.013	0.307	0.036	0.087
SEU	0.001	0.000	0.000	0.001	0.000	0.000	0.000	**0.002**	-0.049	-0.216	0.409	0.427	0.195	0.189	0.122	**0.595**	**0.512**	0.063	0.062	0.145	-0.041
REU	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	**0.002**	0.187	-0.022	-0.019	0.528	-0.045	-0.059	-0.178	-0.121	-0.219	-0.043	0.286	-0.023
LEU	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	**0.001**	-0.086	0.052	-0.061	0.080	-0.009	-0.103	-0.150	-0.127	0.090	0.113	0.023
XEU	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	**0.001**	0.258	0.049	0.017	0.375	0.159	0.224	-0.046	-0.067	0.051	-0.009
SJP	0.001	0.000	0.000	0.000	0.000	0.000	0.000	0.001	0.000	0.000	0.000	**0.002**	-0.189	0.115	0.066	**0.466**	**0.447**	0.145	0.201	0.412	0.171
RJP	0.000	0.002	-0.001	0.001	0.000	-0.001	0.000	0.001	0.005	0.000	0.000	-0.002	**0.069**	0.110	-0.140	-0.120	-0.057	-0.036	-0.075	0.088	0.021
LJP	0.001	0.000	0.000	0.000	-0.001	0.000	0.000	0.001	0.000	0.000	0.000	0.001	0.003	**0.012**	-0.165	0.034	-0.045	0.067	0.010	0.053	-0.075
XJP	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	-0.001	-0.001	**0.001**	-0.038	0.109	-0.159	0.067	-0.154	-0.001
SIE	0.025	-0.007	0.001	0.021	-0.003	-0.002	-0.003	0.029	-0.007	-0.002	0.004	0.021	-0.031	0.004	-0.001	**0.976**	0.699	0.199	-0.059	0.168	0.024
BIE	0.022	-0.004	-0.003	0.020	-0.002	-0.005	-0.005	0.025	-0.005	-0.004	0.005	0.021	-0.015	-0.005	0.004	0.686	**0.988**	0.204	-0.150	-0.040	-0.017
CPIUS	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	**0.000**	-0.091	0.048	-0.064
WSUS	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	**0.000**	0.293	0.305
GDPUS	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	**0.000**	0.481
PSBUS	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	**0.000**

Figure 3.7 Covariance/correlation matrices for the model of Figure 3.6

4. ASSET RETURN AND LIABILITY SCENARIO SIMULATION AND CALIBRATION

4.1. *Simulation and Calibration*

4.1.1. The system model discussed in the previous section is a vector nonlinear second order autoregressive model with a monthly timestep. Given initial values of its state variables, it may be simulated without stochastic innovations as a discrete time deterministic dynamical system defining the mean paths of the state variables. The nonlinear dynamics of this deterministic system may be exceedingly complex and the system may rapidly explode or die to zero values of some variables for certain configurations of the (significant) estimated parameters. Graphical emulation of the central tendencies of the historical path by this deterministic system (see Figure 4.1) is a necessary condition for the generation of realistic scenarios – alternative histories – by Monte Carlo simulation of the stochastic dynamical system. Monte Carlo simulation of this nonlinear vector stochastic difference equation is effected by Euler (first order) stochastic simulation of the independent Gaussian or Student t disturbances which are correlated through the estimated Cholesky factor of the contemporaneous covariance matrix. The implication is that a limited number of estimated parameters – both coefficients and volatilities – may need adjustment to make both the deterministic and corresponding stochastic systems graphically match history (in-sample). Since the impacts of parameter changes is complex due to the nonlinearity of the system, this is not an easy task. Nevertheless, intuitions can be developed to make the achievement of reasonably accurate calibrations tractable and we have developed a prototype graphical interface tool *stochgen 3.0* (Dempster *et al.*, 2002) to aid the process graphically. Ideally, the calibration process itself should be formalised as a nonlinear optimisation problem for some out-of-sample prediction error criterion and we are currently working on limited versions of this. However, the development of appropriate prediction criteria is itself a challenge, to say nothing of the fact that the parameter optimisation problem involving an out-of-sample prediction error criterion is a nonconvex optimisation problem of at least the difficulty of the dynamic stochastic optimisation problems we wish to solve. As previously noted we have therefore made considerable use of graphics.

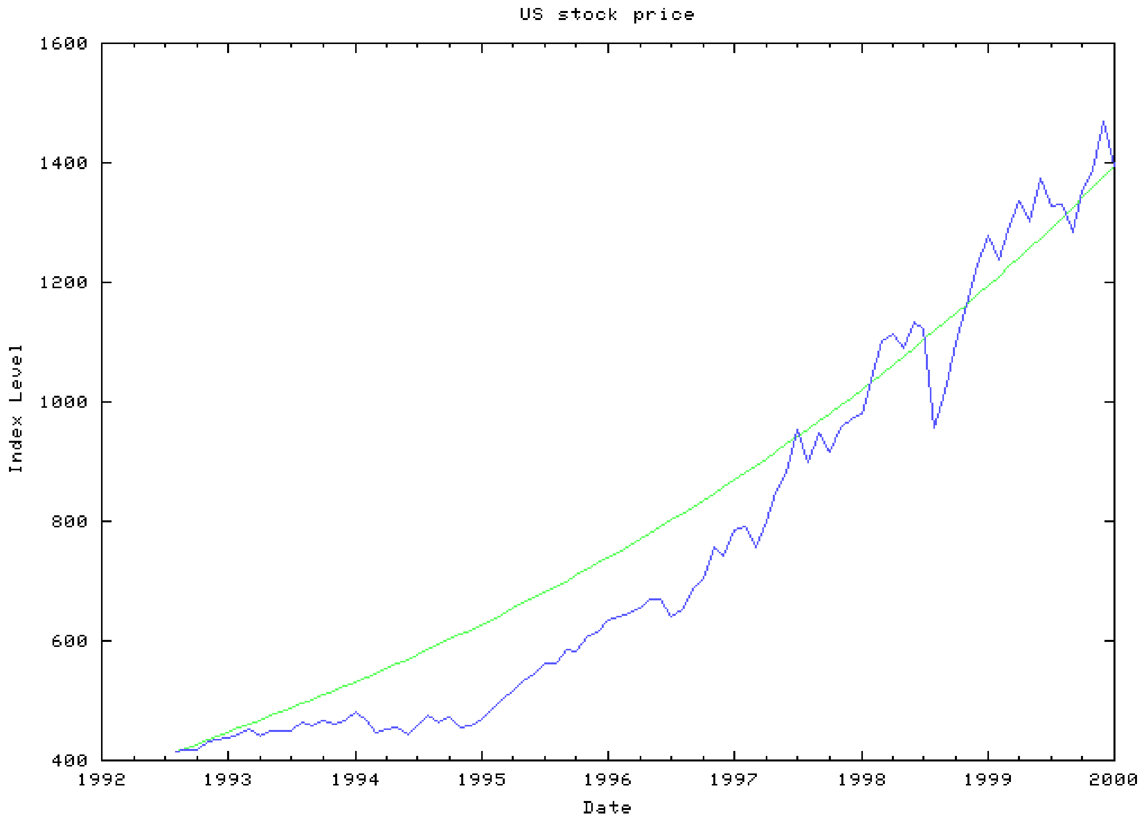

Figure 4.1a In-sample drift with US history – stock index

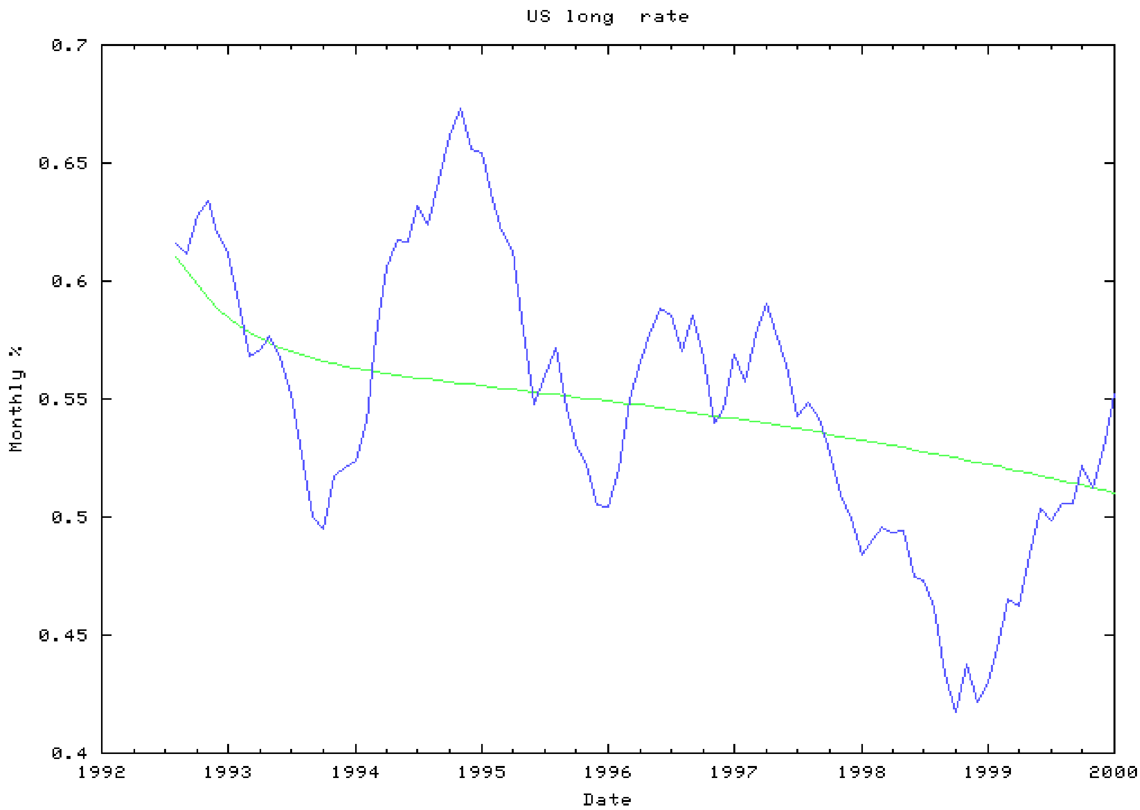

Figure 4.1b In-sample drift with US history – long rate

4.1.2. Figure 4.1 shows a typical graphical result of a calibrated deterministic simulation of the nonlinear system in the estimation period (in-sample). Figure 4.2 shows the corresponding in-sample scenario generation where one is looking for scenario paths with similar properties to the historical path. Similar scenarios may be generated out-of-sample. For calibration purposes however the 0, 25, 50, 75 and 100 percent scenario values in each out-of-sample period, as shown in Figure 4.3, are more valuable. The US stock index plot in the figure shows the desirable calibration in which out-of-sample the historical path is centred in the 50 percent inter-quantile range of the scenario state distributions over time. The US long rate plot shows the less desirable result in which the historical path is captured by the scenario distributions, but is probabilistically over-predicted. As noted above, in- or out-of-sample calibration of *all* variables is difficult and while the weaker criterion may always be met out-of-sample by calibration, in our experience the stronger criterion is usually only met for about 50% of the state variables in a calibration.

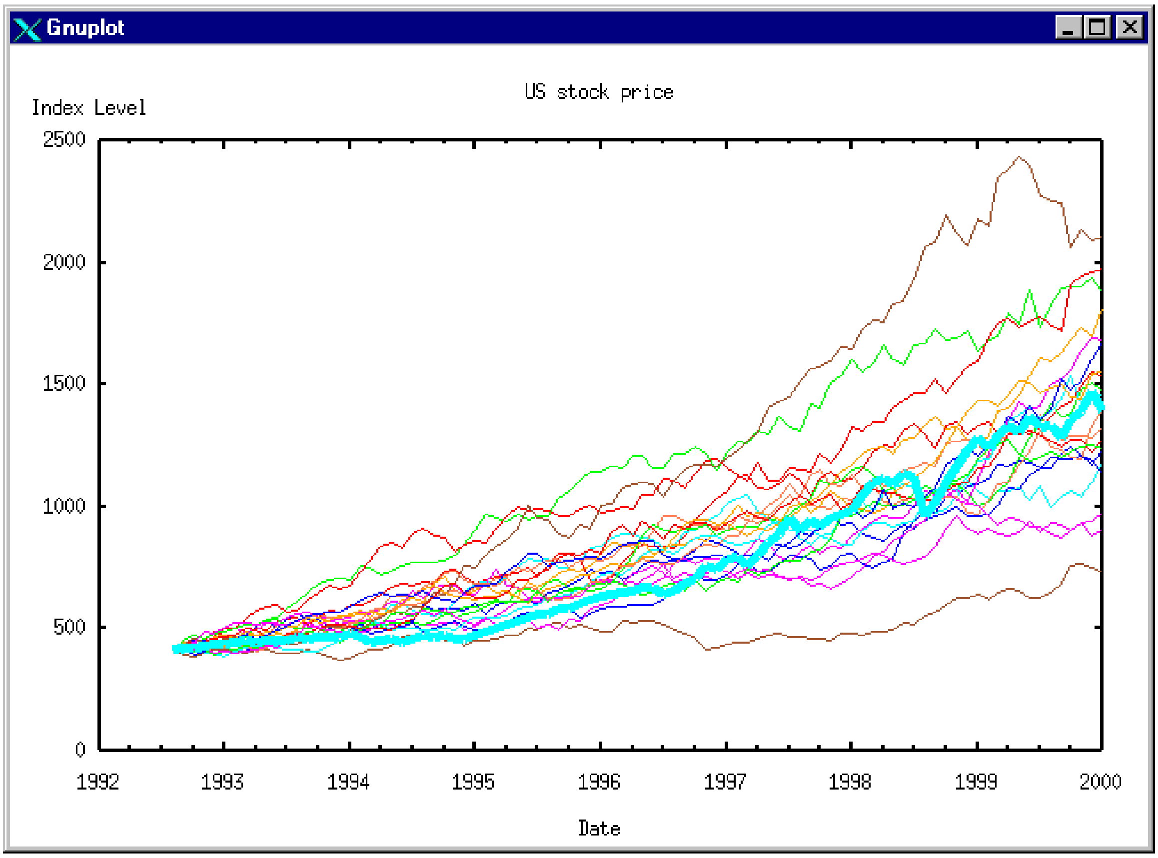

Figure 4.2a In-sample scenarios with US history – stock index

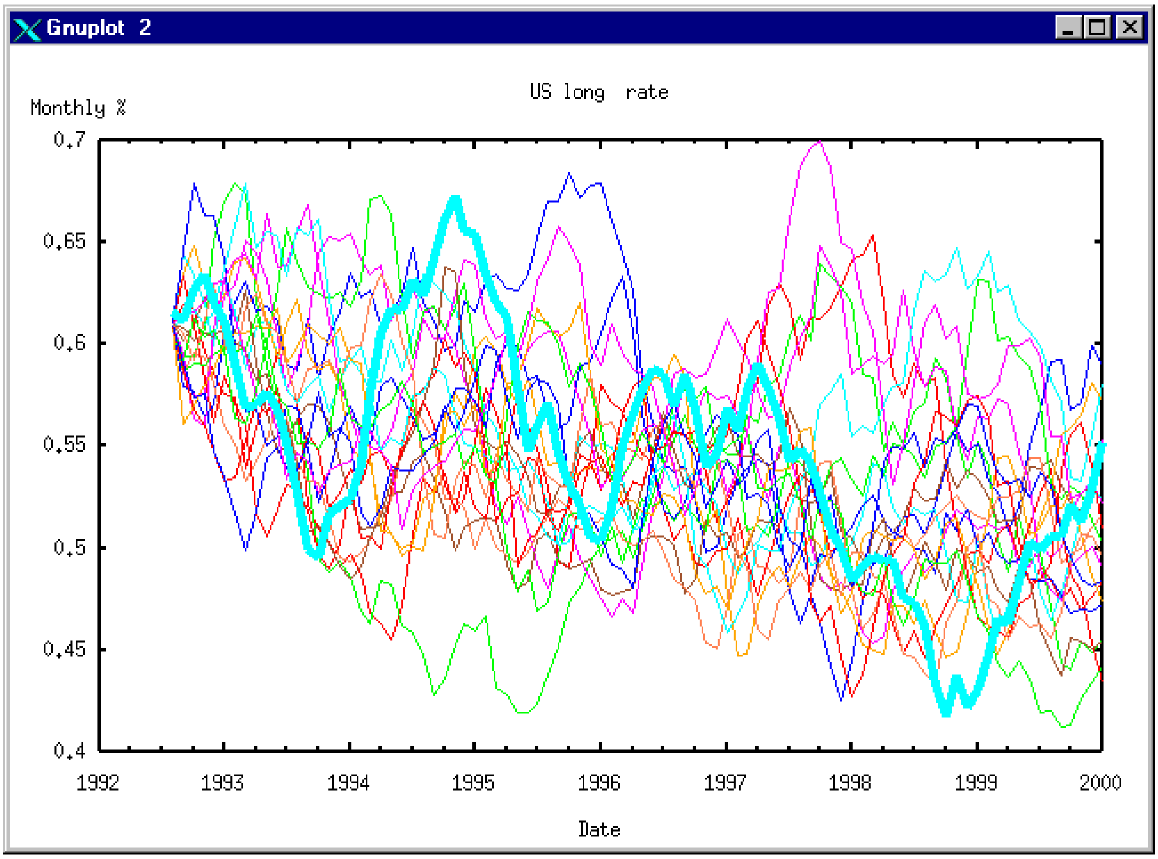

Figure 4.2b In-sample scenarios with US history – long rate

Calibration Period: 77-90

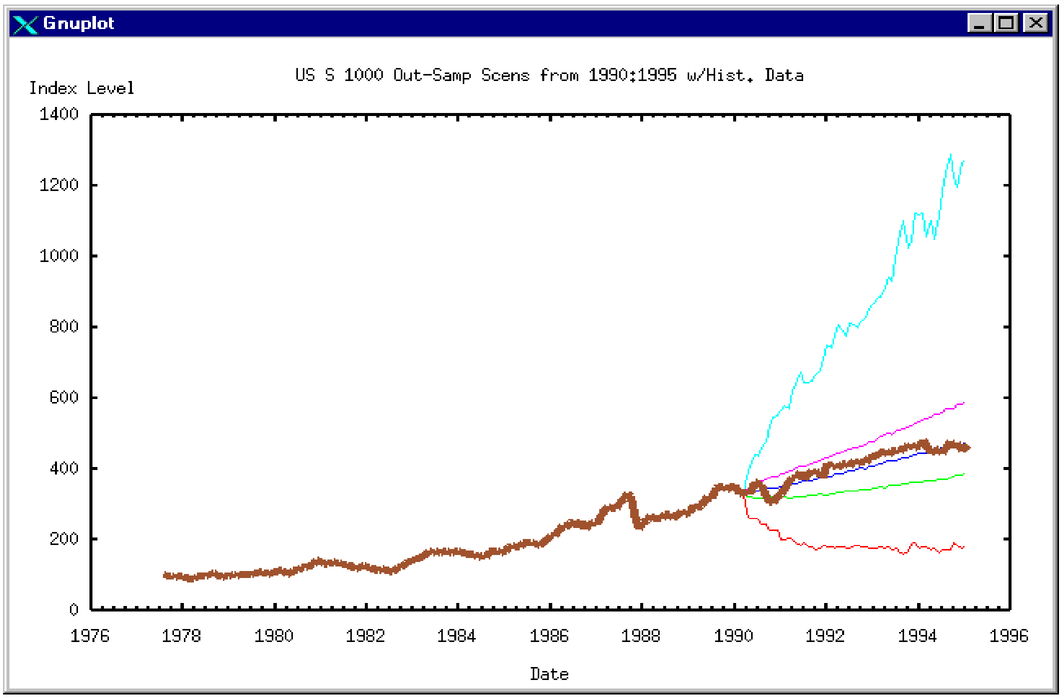

Figure 4.3a Out-of-sample simulation quantiles with US history – stock index

Simulation Period: 90-95

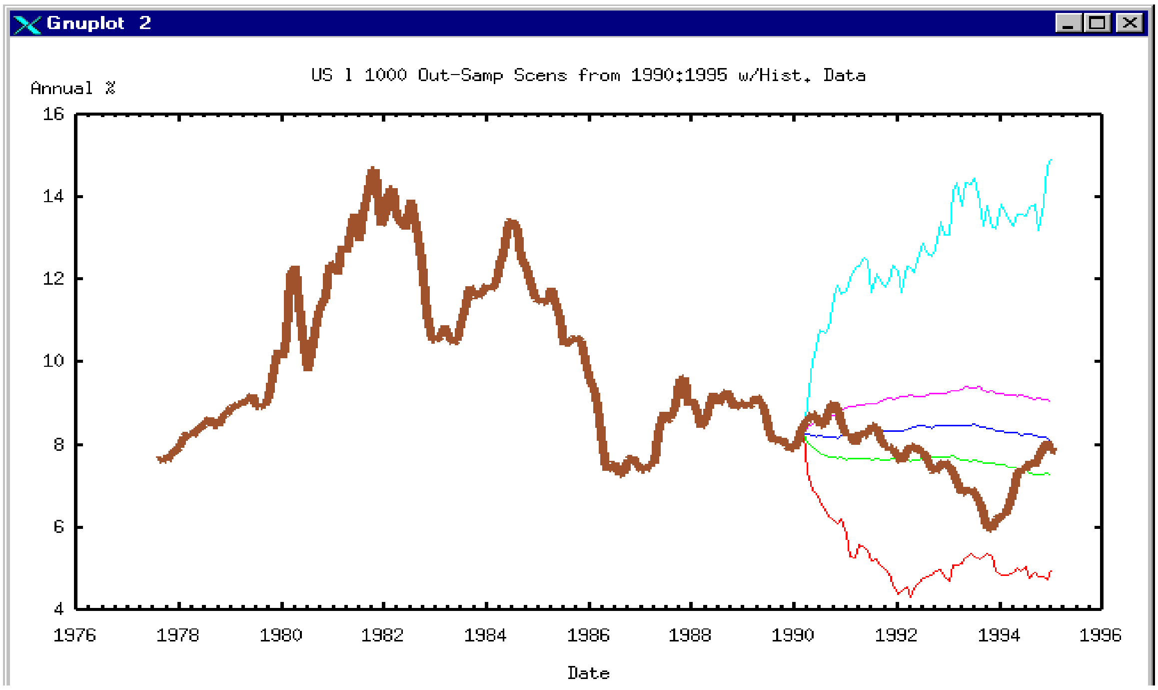

Figure 4.3b Out-of-sample simulation quantiles with US history – long rate

4.1.3. Another approach to econometric model calibration is to linearise a nonlinear system to obtain a *vector autoregressive* (VAR) system which is *stable* in the state variable returns, so that the deterministic system converges to steady state returns and shocks to the corresponding stochastic system are nonpersistent. Stability analysis for such a system is more easily conducted by appropriate eigenvalue analysis of the explanatory variable coefficient matrices – the leading eigenvalue (root) must be less than one in modulus. For given data the feasibility of fitting such a model may be checked by (autoregressive) impulse response analysis (Garratt *et al.*, 2000; Hamilton, 1994) and testing on our full model data to August 2002 has been affirmative. The VAR approach can be extended to an adaptive *error-correcting* VAR model (Boenders *et al.,* 1998; Pesaran & Schuerman, 2001) on which we are currently engaged and will be reported elsewhere (Arbeleche, 2002).

4.1.4. Finally, treating the process generating the historical data as *stationary* with independent increments – an unrealistic assumption – we may alternatively conduct historical simulation by resampling from the empirical marginal distributions of state variable returns constructed from the historical paths over the in-sample period.

4.1.5. All these options have been evaluated and we report dynamic stochastic optimisation backtest results for all three approaches to scenario generation for our dynamic ALM problem in §7.

Calibration Period: 77-90

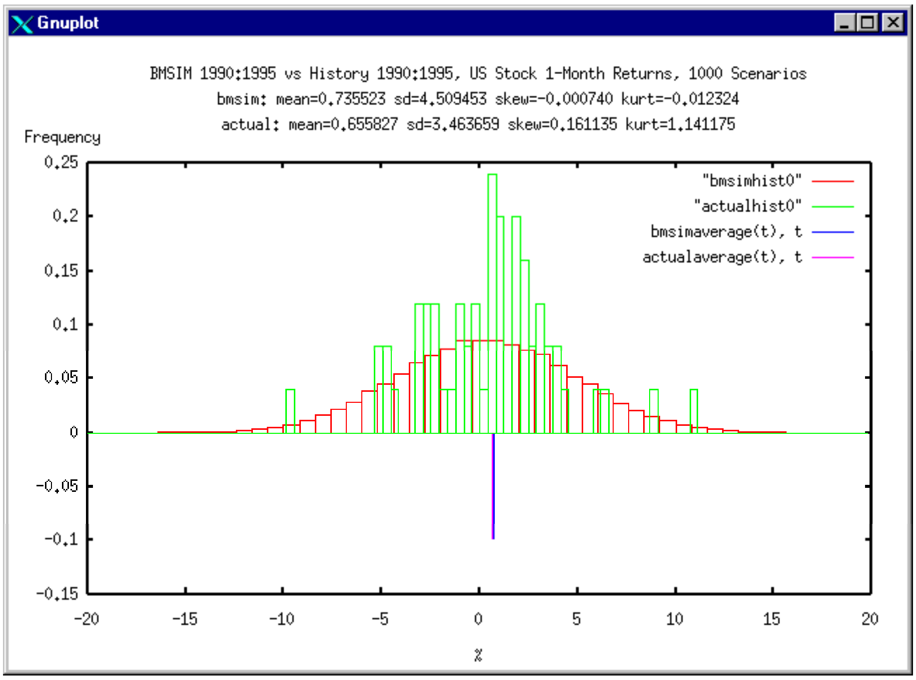

Figure 4.4a Comparison of 1-month returns with US history – stock index

Simulation Period: 90-95

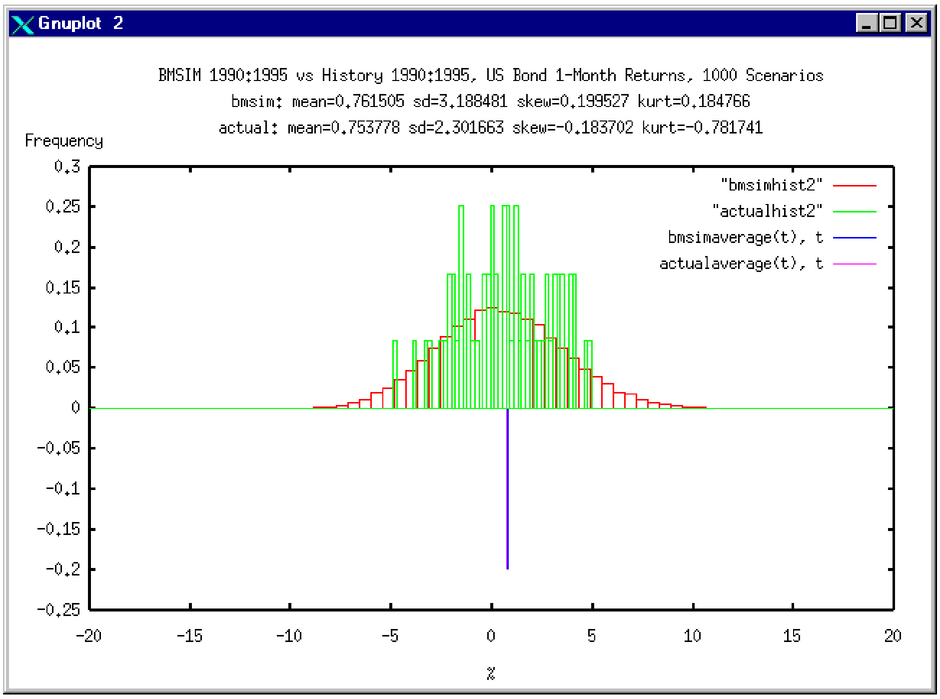

Figure 4.4b Comparison of 1-month returns with US history – long bond

Calibration Period: 92-00

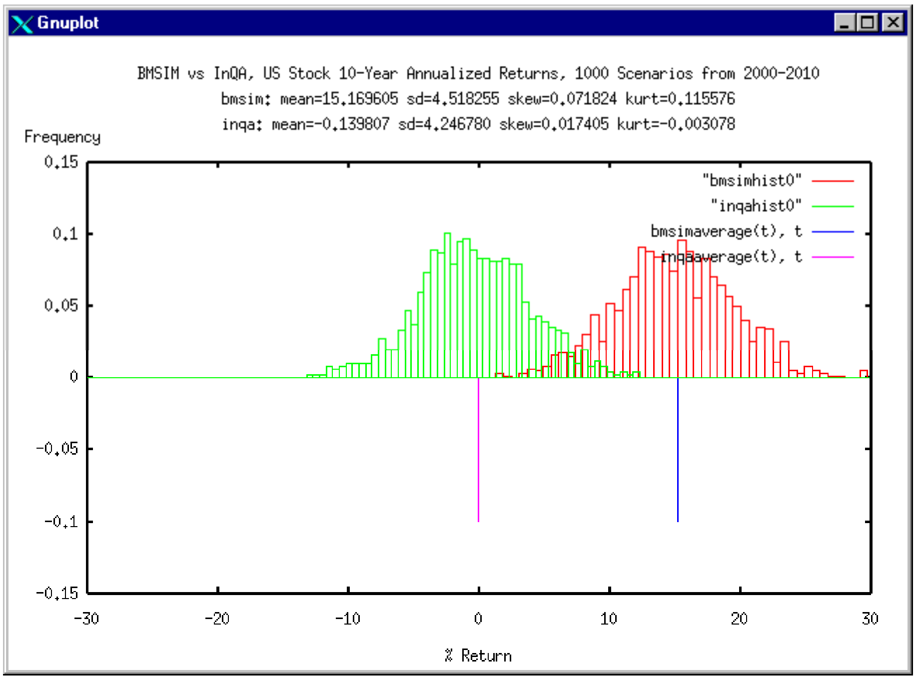

Figure 4.5a Comparison of 10-year annualized US returns with InQA – stock index

Simulation Period: 00-10

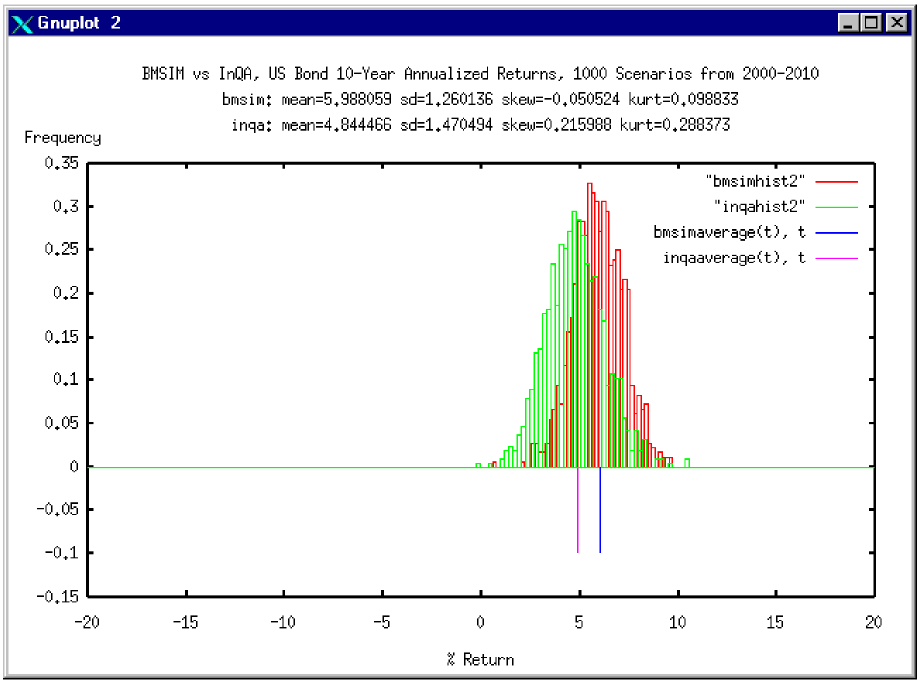

Figure 4.5a Comparison of 10-year annualized US returns with InQA – long bond*Comparative Scenario Return Distribution Evaluation*

Out-of-sample scenario marginal return distributions from calibrated system models were evaluated in two ways: against the empirical marginal return distribution generated by the out-of-sample historical path (Figure 4.4) and against an alternative scenario generation system (Figure 4.5). The 10 year out-of-sample annual return distributions in (the representative) Figure 4.5 were generated by the capital markets model and the market-neutral version of InQA's simulator based on the Wilkie global model (Wilkie, 2000). These comparative results were judged to be more than acceptable.

4.3. *Suppression of Sampling Error*

Since we must always use a finite sample of scenarios, there will always be sampling error in the generation of scenario return state distributions relative to the calibrated estimated system model. This can lead to serious errors and spurious arbitrages in subsequent portfolio optimisation. These can however be suppressed by ensuring that the sample marginal return distributions corresponding to all generated scenarios at a specific point in time have two moments matched to those of the theoretical model underlying the simulations (Høyland & Wallace, 2001; Høyland *et al.,* 2001). This can be posed in terms of matching the moments of the sampled innovations with their theoretical –here independent standard normal or student t – distributions. The first sample moments are easily set to zero by translation and the unit second moments can be matched in terms of a nonlinear programme which can be solved by sequential quadratic programming using the SNOPT sequential quadratic programming software (Villaverde, 2002).

4.4. *Liability Modelling and Simulation*

4.4.1. A proprietary stochastic Markov chain model for defined benefit pension fund liabilities has been developed which currently assumes (unrealistically) that liabilities and fund return performance and macroeconomic variables such as CPI and the wages and salaries index are independent. Nevertheless, formidable calibration problems for the liability model remain due to lack of historical data.

4.4.2. For defined contribution pension funds similar interdependence between lagged fund performance and participation rates is a reality. In principle this can be handled (Dempster, 1988), but is again difficult to specify and calibrate.

4.4.3. Tax liabilities for funds in the jurisdiction of the fund manager are particularly simple – a one percent proportional transaction cost.

4.4.4. If complex liability models (including more complex tax liabilities) can be simulated – possibly together with asset returns and macroeconomic variables – to result in a net liability cash flow process, no difficulties arise in the optimisation model (see e.g. Consigli & Dempster, 1998). In this paper however we concentrate on the newer – previously unsolved – problem of incorporating the guarantee liabilities of defined contribution pension plans into scenario based stochastic optimisation models (see §6.4).

4.5. *Scenario Tree Generation*

4.5.1. As mentioned in §2.2, in order to mirror reality dynamic stochastic optimisation models for strategic DFA problems must face alternative scenario uncertainty at each decision point in the model – e.g. at each forward portfolio rebalance. Otherwise, the model decisions incorporate future knowledge along scenarios– hardly possible in the real world of finance! The distinction is between the so-called *flat* out-of-sample scenarios of Figure 4.6 and a scenario *tree,* an example of which is shown schematically in Figure 4.7. Each path from the root to a leaf node in the latter scheme *represents* a scenario and the nodes represent decision points – the root node represents the initial *implemented* decision (e.g. initial portfolio balance). Subsequent nodes represent *forward* 'what-if' decisions facing the uncertainty represented by all scenarios emanating from that node.

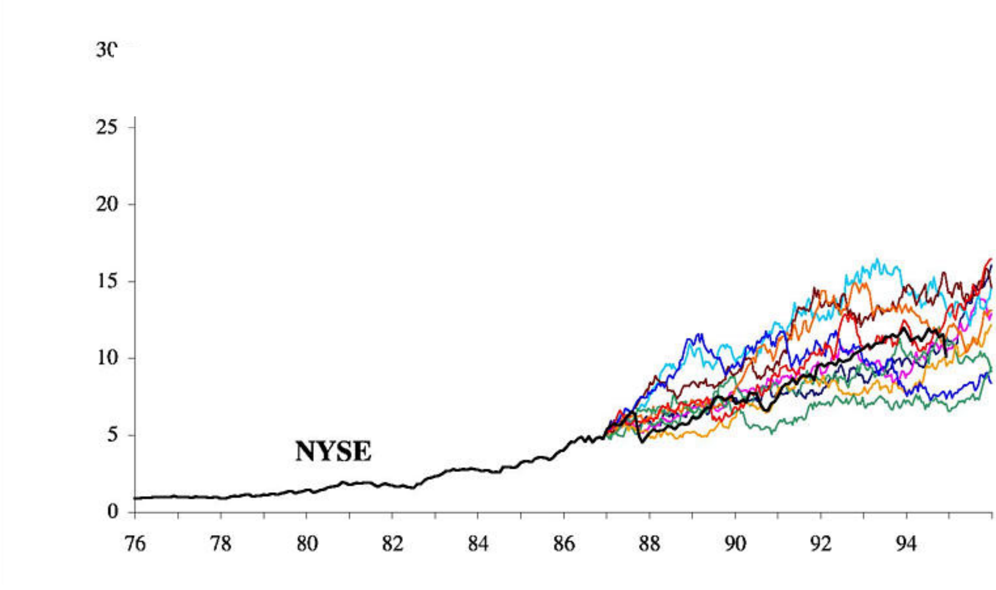

Figure 4.6 Out-of-sample flat scenario generation

4.5.2. Note that the Monte Carlo simulation of scenarios corresponding to a given scheme is a nontrivial matter requiring generic software to handle a complex simulator such as is needed for the Pioneer model. We have used the generic *stochgen 2.3* software of the STOCHASTICS™ toolchain for dynamic stochastic optimisation (Dempster *et al*, 2002) and its variants tailored for Pioneer.

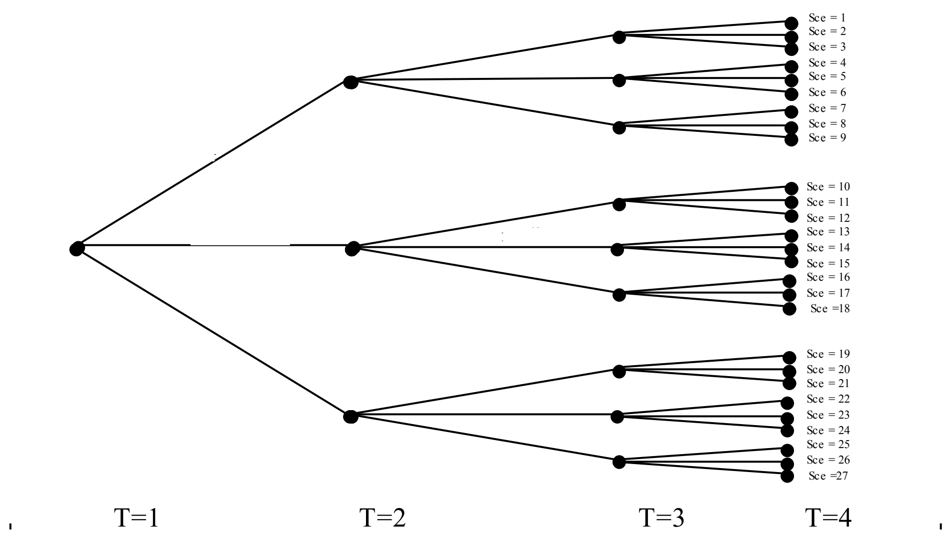

Figure 4.7 Schematic out-of-sample scenario tree branching structure with uniform branching factor 3

4.5.3. In this software the input tree structure is represented for a symmetric *balanced* scenario tree by a product of branching factors, e.g. 3.3.3 or 3^3 for the scenario tree of Figure 4.7, or by a scenario or nodal partition matrix for asymmetric trees as shown in Figure 4.8. The *scenario partition matrix* (Lane & Hutchinson, 1980) corresponds to the discrete scenario information partition inherent in the tree structure at each decision point while the *nodal partition matrix* (used in *stochgen 2.3*) denotes the node through which each scenario passes at each decision point and is useful for decomposition-based optimisers.

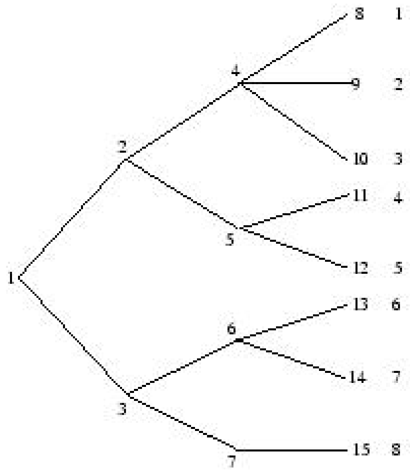

$$L = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 2 \\ 1 & 1 & 1 & 3 \\ 1 & 1 & 4 & 4 \\ 1 & 1 & 4 & 5 \\ 1 & 6 & 6 & 6 \\ 1 & 6 & 6 & 7 \\ 1 & 6 & 8 & 8 \end{bmatrix} \qquad M = \begin{bmatrix} 1 & 2 & 4 & 8 \\ 1 & 2 & 4 & 9 \\ 1 & 2 & 4 & 10 \\ 1 & 2 & 5 & 11 \\ 1 & 2 & 5 & 12 \\ 1 & 3 & 6 & 13 \\ 1 & 3 & 6 & 14 \\ 1 & 3 & 7 & 15 \end{bmatrix}$$

Figure 4.8 Example of a scenario tree with corresponding scenario and nodal partition matrices

4.5.4. The *stochgen* software must handle at each node multiple conditional stochastic simulations of versions of the asset return model initialised by the data at the node and two previous timesteps (months) along the scenario path. Notice that the simulation time step (a month) is much shorter than the decision point frequency (for forward portfolio rebalancing: quarterly, semi annually or annually), *cf.* Dempster *et al.* (2000).

4.5.5. In the reported project backtests we used balanced scenario trees with high initial branching (see §7.2).

4.5.6. A number of variants of the BMSIM stochastic simulator for the Pioneer model have so far been written in C^{++}/C, but in the *stochgen 3.1* software currently under development these variants are specified as extensions or restrictions of a full model. Similarly VARSIM and VARSIM 2 are simulators for variants of the VAR linearisation of the asset return model and HSIM performs the historical bootstrap simulation described above in §4.1.

4.5.7. Obtaining bond returns *in* a currency area is somewhat subtle since they must be derived from bond yields. A representative derivation is given in Appendix B. Handling a complex external stochastic simulator is just one function of the variants of the *stochgen* software and we will return to its other functions in §6 after describing in the next section the strategic ALM dynamic stochastic optimisation models used in our project.

5. OPTIMAL DYNAMIC ASSET LIABILITY MANAGEMENT

5.1. *CALM Problem Formulation*

5.1.1. The dynamic ALM model used in the Pioneer project is a variant of the CALM (computer-aided asset liability management) model (Dempster, 1993) used previously in other projects (Consigli & Dempster, 1998; Hicks Pedrón, 1998). Here we describe the main features of the model. A precise mathematical description is given in Appendix C.

5.1.2. We focus in this paper on what is normally called *strategic asset allocation* which is concerned with allocation across broad asset classes such as equity and bonds of a given country. The problem is as follows:

Given a set of assets, a fixed planning horizon and a set of rebalance dates, find the trading strategy that maximizes the risk adjusted wealth accumulation process subject to the constraints.

As noted in §4.4, defined contribution pension plans or other complex liabilities (such as insurance or reinsurance claims) may be added to the basic model as a stochastic net cashflow stream (see e.g. Consigli and Dempster, 1998).

5.1.3. In the model description given below we begin with a discussion of alternative utility functions (fund risk tolerances) (§5.2) and continue on to treat the specification of risk management objectives through the problem objective function (§5.3) and then the constraints (§5.4). The last two sections discuss respectively the optimal setting of benchmark portfolios (§5.5) and the specification of probabilistic *value at risk* (VaR) constraints for the model connected with defined contribution guarantee liabilities (§5.6).

5.1.4. We consider a discrete time and space setting. It is assumed that the fund operates from the view point of one currency which we call the *home currency*. Unless otherwise mentioned all quantities are assumed to be in the local currency. There are T+1 times (the first T are decision points) indexed by t=1,...,T+1, where T+1 corresponds to the planning *horizon* at which no decisions are made. Uncertainty is represented by a finite set of time evolutions of states of the world, or *scenarios,* denoted by Ω. The probability $p(\omega)$ of scenario ω in Ω is here always the reciprocal of the number of scenarios since these scenarios are being generated by Monte Carlo simulation as discussed in the previous section.

5.1.5. Assets take the form of *equity*, *bonds* and *cash*. Let I denote the set of all equity and bond assets and K denote the set of cash assets. The fund begins with an *initial endowment* of equity and bonds given by $\{x_i: i \in I\}$ and of cash in the home currency given by w_1. The fund trades in the assets at t=1,...,T, i.e. at all times except for at the planning horizon.

5.1.6. A *trading* strategy is given by $\theta_{ikt}(\omega):=(x_{it}(\omega),x_{it}^{+}(\omega),x_{it}^{-}(\omega),z_{kt}^{+}(\omega),z_{kt}^{-}(\omega))$ for i in I, k in K, t=1,..., T, ω in Ω, where:

- $x_{it}(\omega)$ denotes the amount *held* of asset i between time t and time t+1 in state ω.

- $x_{it}^{+}(\omega)/x_{it}^{-}(\omega)$ denotes the amount *bought/sold* of asset i at time t in state ω. The introduction of the buy/sell variables is used to account for proportional transaction costs on buying and selling equity and bond assets. Denote by f and g respectively the *proportional transaction cost* of buying or selling an equity or bond asset.. For example, a 1% proportional transaction cost on buying and selling an equity or bond asset corresponds to f=1.01 and g=0.99.

- $z_{kt}^{+}(\omega)/z_{kt}^{-}(\omega)$ denotes the amount of cash *lent/borrowed* in asset k between time t and time t+1 in state ω. The positions in cash are split into long and short components to account for different rates of borrowing and lending. We assume that cash lent and borrowed at time t in any currency is automatically converted back to the home currency at time t+1.

- The *asset returns* are given by $\{v_{it}(\omega), (r_{kt}^{+}(\omega),r_{kt}^{-}(\omega))$: i in I, k in K, t=2,...,T+1, ω in $\Omega\}$ where:

- $v_{it}(\omega)$ denotes the net return on asset i between time t-1 and time t in state ω.

- $r_{kt}^{+}(\omega)/r_{kt}^{-}(\omega)$ denotes the net return on *lending*/*borrowing* asset k between time t-1 and time t in state ω.

5.1.7. The *exchange rates* are given by {($p_{it}(\omega)$,$p_{kt}(\omega)$): i in I, k in K, t=1,...,T+1, ω in Ω} where:

- $p_{it}(\omega)$ denotes the exchange rate of asset i at time t in state ω expressed as home currency/local currency.

- $p_{kt}(\omega)$ denotes the exchange rate of asset k at time t in state ω expressed as home currency/local currency.

5.1.8. The fund may face cash inflows and outflows given by {($q_t^{+}(\omega)$,$q_t^{-}(\omega)$): t=2,...,T, ω in Ω} where:

- $q_t^{+}(\omega)/q_t^{-}(\omega)$ denotes the *cash flow in/out* at time t in state ω.

5.1.9. A trading strategy θ results in a *wealth before rebalancing* of $w_t^{\theta}(\omega)$ for t=2,...,T+1 and ω ∈Ω, and a *wealth after rebalancing* of $W_t^{\theta}(\omega)$ for t=1,…,T and ω ∈Ω.

5.1.10. Subject to the constraint structure, the fund acts by choosing the trading strategy which maximizes *the* (von Neumann-Morgenstern) *expected utility* of the wealth process which is assumed to take the form

$$E[U(\mathbf{w}_2^{\theta},...,\mathbf{w}_{T+1}^{\theta})] = \sum_{\omega\in\Omega} p(\omega)\sum_{t=2}^{T+1} u_t(w_t^{\theta}(\omega)) .$$

Alternative *period utility functions* u_t are discussed in the next section.

5.2. Utility Functions

5.2.1. The functional U is used to define the risk preferences of the fund over the wealth process in such a way that $E[U(\mathbf{w}_2^{\theta_1},...,\mathbf{w}_{T+1}^{\theta_1})] > E[U(\mathbf{w}_2^{\theta_2},...,\mathbf{w}_{T+1}^{\theta_2})]$ if and only if the wealth process generated by θ_1 is strictly preferred to the wealth process generated by θ_2. Thus a clearly desirable property of U is that it be strictly increasing. Another desirable property of U is that it be concave. If U is concave $E[U(\mathbf{w}_2^{\theta},...,\mathbf{w}_{T+1}^{\theta})] \le U(E[\mathbf{w}_2^{\theta}],...,E[\mathbf{w}_{T+1}^{\theta}])$. The interpretation is that the utility of having the certain quantities $E[\mathbf{w}_2^{\theta}],...,E[\mathbf{w}_{T+1}^{\theta}]$ is preferred to the expected utility of having the uncertain quantities $\mathbf{w}_2^{\theta},...,\mathbf{w}_{T+1}^{\theta}$. Thus if U is concave the fund is said to be *risk-averse* and if it is linear it is said to be *risk-neutral*. (If U is convex then it is said to be *risk-loving* or *risk-seeking*.) Since U is a linear combination of the u_t, U will be strictly increasing and concave if they are.

5.2.2. As noted in §2.1 the utility functional is used here to represent the general attitude to risk of the fund's participants over a specified fund horizon. Short horizon funds are likely to attract more risk averse participants than very long horizon funds whose long term participants can afford to tolerate more risk in the short run. Even for such problems however the fund manager will likely wish to mitigate the long term participants' risk tolerances in the short run in the interest of maintaining competitive participation rates. In any event, choice of a sequence of period utility functions can be used to shape the evolution of the wealth process over the scenarios in the scenario tree of the problem. Appropriate tree size and branching structure – together with variance reduction (§4.3) – can be used to ensure that these distributional properties resulting from the implemented decisions continue to hold against sufficiently large samples of further flat scenarios not included in the problem scenario tree – a prerequisite for good out-of-sample performance (see §7.2).

5.2.3. We consider the following period utility functions:

1. *Exponential* (CARA): $u(w) = -e^{-aw}$ $\quad a > 0$

2. *Power* (CRRA): $u(w) = \dfrac{w^{1-a}}{1-a}$ $\quad a > 0$

3. *Downside-quadratic*: $u(w) = (1-a)w - a(w-\tilde{w})_{-}^{2}$ $\quad 0 \le a \le 1, 0 \le \tilde{w} \le \infty$.

5.2.4. Note that *log* utility given by $u(w) = \log(w)$ is a limiting case of power utility as $a \to 1$. The $\tilde{w}$ parameter that appears in the downside-quadratic utility function denotes a *target wealth*. Note that this utility function reduces to *linear* (risk-neutral) utility given by $u(w) = w$ for a:=0.

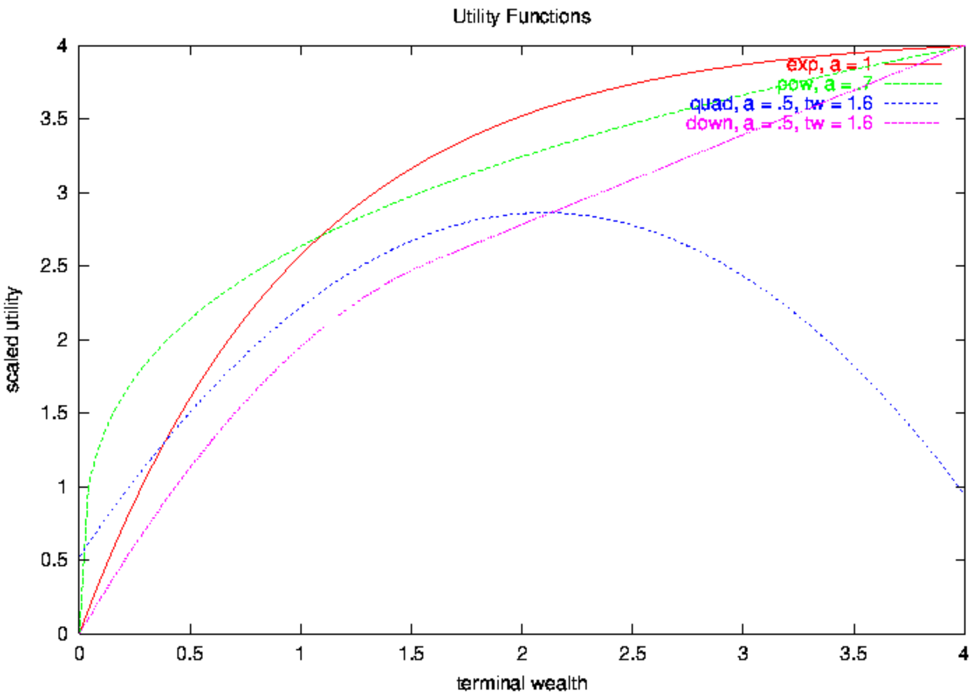

Figure 4.9 Scaled risk averse utility functions

5.2.5. The exponential utility function is also referred to as the constant *absolute risk aversion* (CARA) utility function because its Arrow-Pratt *absolute measure* of risk aversion defined by $-u''(w)/u'(w)$ is equal to the constant a. The power utility function is also referred to as the *constant relative risk aversion* (CRRA) utility function because the Arrow-Pratt *relative measure* of risk aversion defined by $-wu''(w)/u'(w)$ is equal to the constant a. The downside-quadratic utility function, similar to the *mean-downside-variance* or *mean-semi-variance* utility function except that it has $\tilde{w}$ in place of E[**w**], aims to maximize wealth and at the same time penalize downside deviations of the wealth from the target. This is illustrated in Figure 4.9 which depicts the different amounts of risk aversion implied by the curvature of the utility functions – for a fixed slope the greater the curvature the greater the aversion to risk. Of particular interest is the curvature for wealth levels less than the initial wealth (1) or the target wealth ($\tilde{w}$).

5.2.6. Table 4.2 gives the Arrow-Pratt absolute measure of risk aversion for each utility function considered above and used in our models.

Exponential	a
Power	a/w
Downside-Quadratic	$2a/[(1-a)-2a(w-\tilde{w})]$

Table 4.2 Arrow-Pratt Absolute Measure of Risk Aversion

Kallberg and Ziemba (1983) have shown in the one period case that utility functions with similar Arrow-Pratt absolute measures of risk aversion result in similar optimal portfolios.

5.3. Risk Management Objectives

As noted above, in principle different attitudes to downside risk in fund wealth may be imposed at *each* decision point through the *additively separable* utility U which is a sum of different period utility functions u_t, t=2,…,T+1, or may be of a common form with different period-specific values of its parameters. Adjustment of these parameter values allows the shaping of the fund wealth distribution across scenarios at a decision point as we shall see in more detail in §6. In practice however a common specification of period utility is usually used.

5.4. Basic, Diversification and Liquidity Constraints

5.4.1. The basic constraints of the dynamic CALM model (*cf.* Consigli and Dempster, 1998) detailed in Appendix C are:

- *Cash balance constraints.* These are the first set of constraints of the model referring respectively to period 1 and the remaining periods before the horizon.

- *Inventory balance constraints*. These are the second set of constraints and involve *buy* (+), *sell* (-), and *hold* variables for each asset (and more generally liability, with buy and sell replaced by incur and discharge). This approach, due to Bradley & Crane (1972), allows (with double subscripting) all possible tax and business modelling structures to be incorporated in constraints (see e.g. Cariño et al., 1994).

- *Current wealth constraints*. The third set of constraints involves the two wealth variables: beginning of period *wealth before rebalancing* (w) from the previous period and beginning of period *wealth* (W) after a possible cash infusion from borrowing, or an outflow from the costs of portfolio rebalancing and possible debt reduction, i.e. *after rebalancing*.

5.4.2. The remaining constraint structures required will likely differ from fund to fund. Possible constraints include:

- *Solvency constraints*. These constrain the net wealth of the fund generated by the trading strategy θ to be non-negative (or greater than a suitable regulatory constant) at each time, i.e. $w_t^\theta(\omega) \geq 0$ for t=2,…,T+1 and ω in Ω.

- *Cash borrowing limits*. These limit the amount the trading strategy can borrow in cash and take the form: $p_{kt}(\omega) z_{kt}^-(\omega) \leq \bar{z}_k$ for k in K, t=1,…,T and ω in Ω, where recall p_k denotes the appropriate exchange rate.

- *Short sale constraints*. These limit the amount the trading strategy can short the equity and bond assets and take the form: $p_{it}(\omega) x_{it}(\omega) \geq \bar{x}_i$ for i in I, t=1,…,T and ω in Ω.

- *Position limits*. These limit the amount invested in an asset to be less than some proportion $\phi < 1$ of the fund wealth and take the form:

$$p_{it}(\omega) x_{it}(\omega) \leq \phi_i W_t^\theta(\omega)$$
$$p_{kt}(\omega)(z_{kt}^+(\omega) - z_{kt}^-(\omega)) \leq \phi_k W_t^\theta(\omega)$$

for i in I, k in K, t=1,…,T and ω in Ω.

- *Turnover (liquidity) constraints*. These limit the approximate change in the fraction of total wealth invested in some equity or bond asset i from one time to the next to be less than some proportion of the fund wealth $\alpha_i < 1$ and take the form:

$$| p_{it}(\omega)x_{it}(\omega) - p_{it-1}(\omega)x_{it-1}(\omega) | \leq \alpha_i W^0_{t-1}(\omega)$$

for i in I, k in K, t=1,…,T and ω in Ω. They are imposed on large funds primarily from market liquidity considerations which are not modelled.

5.4.3. All the above constraints are piecewise linear convex.

5.4.4. For backtesting purposes (see §7) we define the following three types of constraint structures. T1 constraints have no position limits or turnover constraints. T2 constraints have 20% position limits on all assets and no turnover constraints. T3 constraints contain both position limits and turnover constraints as summarized in Table 4.3.

Asset	**Position Limit**	**Turnover Constraint**
US Equity	0.40	0.15
US Bonds	0.40	0.15
UK Equity	0.80	0.15
UK Bonds	0.80	0.15
EU Equity	0.80	0.15
EU Bonds	0.80	0.15
JP Equity	0.15	0.15
JP Bonds	0.15	0.15
EM Equity	0.05	-
EM Bonds	0.05	-
Sum of Cash	0.25	-
US Equity + Bonds	0.50	-
JP Equity + Bonds	0.20	-
EM Equity + Bonds	0.08	-

Table 4.3 Position limits and turnover restrictions by proportion of value

5.4.5. Short selling and borrowing are not allowed in any of these constraint structures. Assuming that the simulated price processes are non-negative, this automatically enforces the solvency constraints.

5.5. Benchmark Portfolio (Fixed Mix) Constraints

5.5.1. A common problem in the management of funds of all types is the setting of realistic benchmarks. This is usually done in an *ad hoc* manner in light of experience. For a given set of asset classes a *benchmark portfolio* whose performance can be used to set a *return benchmark* may be decided *optimally* by applying a further constraint to any variant of the dynamic strategic ALM model so far defined. The corresponding portfolio rebalance (trading) strategy is to rebalance the asset portfolio to the initial optimally determined proportions – i.e. fixed mix – at each trading date (decision point), see Mulvey (1995). Thus assets which have appreciated since the last rebalance will be sold to finance the purchase of depreciating assets to bring their value up to the initial fixed proportion of portfolio value – *buy low and sell high*! – but of course this policy is no protection against generally falling asset values.

5.5.2. Mathematically, the *fixed mix* constraint on asset values held in each scenario ω in Ω at each time period t=1,…,T is given by

$$\Sigma_{i\in I}\lambda_i = 1$$
$$p_{i1}x_{i1}^{+} = \lambda_i\left(w_0 - \tau\right) \qquad i \in I$$
$$p_{it}\left(\omega\right)x_{it}\left(\omega\right) = \lambda_i\left(\sum_{j\in I} p_{jt}\left(\omega\right)x_{jt}\left(\omega\right)\right) \qquad i \in I\ ,$$

where $\lambda_i \geq 0,\ i \in I$, are the initial portfolio *proportions* to be optimally determined, w_0 is initial wealth and τ is an estimate of the transaction costs of the initial portfolio balance. Obviously the imposition of these constraints reduces the terminal wealth achievable in the model relative to the full optimum without such constraints – sometimes *severely* in practice (Hicks Pedrón, 1998) – and hence constitutes a benchmark to beat. Unfortunately, due to the bilinear nature of the constraints applying to the portfolio decisions subsequent to the initial one the resulting optimisation problem becomes *nonconvex* (Dempster *et al.*, 2003), but we shall address its practical solution in §6.

5.6. *Guaranteed Return Constraints*

5.6.1. Of course the return guarantee to an individual investor in a defined contribution pension fund is absolute, given the solvency of the guarantor. In the situation of a banking group such as the fund manager and its parent guarantor this necessitates strategies both to implement the absolute guarantee for individuals and to manage the investment (trading) strategy of the fund so as to ensure meeting the guarantee for *all* participants of the fund with a high probability.

5.6.2. Mathematically, this latter goal can be met by imposing a *probabilistic* constraint of the VaR type on the wealth process at specific trading dates, computing expected shortfall across scenarios which fail to meet the fund guarantee and adding the corresponding penalty terms to period objective functions. For example, at the horizon T+1 or any intermediate date t' this would take the form

$$P\left(\mathbf{w}_{t'} \geq w_{t'}^{*}\right) \geq 1-\alpha,$$

where α: = 0.01 or 0.05, corresponding to respectively 99% or 95% confidence, and $w_{t'}^{*}$ is calculated from the initial wealth and the guaranteed annualised rate *r* as $w_0\left(1+r\right)^{t'}$. However, such scenario-based probabilistic constraints are extremely difficult to implement in that they again convert the convex (deterministic equivalent) large scale optimisation problem to a *nonconvex* one. We will nevertheless describe a practical approximation procedure in the next section, but we leave expected shortfall penalties to future work.

6. PROBLEM GENERATION AND SOLUTION TECHNIQUES

6.1. *Optimisation Problem Generation*

6.1.1. Instantiations of the CALM model and other similar strategic DFA models lead to *very large* deterministic equivalent nonlinear optimisation problems involving perhaps hundreds of thousands of scenarios and millions of variables and constraints. Moreover in a production setting both parameter values and the model itself are constantly changing due to changes of view, objectives and regulations. Mathematical programming *modelling languages* such as AMPL (Fourer *et al.*, 1993) and OPL (ILOG, 2000) have been developed to handle deterministic optimisation models in this regard by specifying the variables, objective and constraints of the problem in an algebraic language in terms of entity sets which is similar to the ordinary mathematical specification of the Pioneer CALM model given in Appendix C. Such systems take as input the model in algebraic form together with specific parameter values and they output a structured file in a *standard* format such as MPS (IBM, 1972) which is readable as input by a wide range of optimisation solvers. These concepts have been extended to large scale dynamic stochastic optimisation problems with the STOCHASTICS™ software (Dempster *et al.*, 2002) and the SMPS standard solver input format (Birge *et al.*, 1986) which have been used for this project. As discussed in §4.4 the *stochgen* subsystem handles the scenario tree generation using routine dynamic stochastic simulation from a standardised tree structure specification – horizon and branching structure –and making use of AMPL (or a new modelling language SAMPL currently under development for *stochgen 3.1*) outputs the optimisation problem for decomposition based techniques – or appropriate pieces of the optimisation problem – in the SMPS or MPS formats to the solver – possibly as it runs. See Dempster & Consigli (1998) and Dempster *et al.* (2002) for more details.

6.2. *Optimal Strategic ALM Algorithms and Software*

6.2.1. A variety of large scale optimisation algorithms have been used to solve variants of the CALM model. For linear and quadratic problems – both linearly constrained – these are simplex, interior point and nested Benders decomposition methods. For general linearly constrained convex and general nonlinear problems both nested Benders decomposition and sequential quadratic programming algorithms have been used.

6.2.2. Simplex and interior point algorithms are well documented (see e.g. Vanderbei, 2002) and the basic reference to nested Benders decomposition is Gassmann (1990), see also Scott (2002). Nested Benders decomposition is a sequential cutting plane technique in which the subproblems at each node of the scenario tree are solved independently for each major iteration until the cuts for each subproblem lead to the solution of the problem. Like interior point methods, the number of major iterations required for convergence by nested Benders decomposition depends more upon the size of the feasible region than on the problem dimensions (size) itself. We have used CPLEX 5.1 for linear and quadratic programming, *solgen 1.2* of the STOCHASTICS™ toolchain for nested Benders decomposition and SNOPT for general nonlinear programming by sequential quadratic programming (see Gill *et al.*, 2002).

6.2.3. For the CALM model of Appendix C and its variants - which are linearly constrained convex problems generally and quadratic problems for the best performing downside quadratic utilities (see §7.3), we usually first solve a quadratic version of a new instantiation with a few thousand scenario tree using CPLEX interior point. For the very large scenario trees corresponding to long horizon multi-portfolio rebalance problems however the *solgen 1.2* implementation of nested Benders decomposition is required since the other techniques must load the full problem into the computer's memory.

6.3. *Optimal Benchmark Portfolio Algorithms*

6.3.1. Due to the bilinear nature of the constraints – in initial portfolio proportions and subsequent portfolio asset positions – which apply to portfolio decisions subsequent to the initial one the fixed mix problem for setting optimal benchmark portfolios is *nonconvex*. However these nonconvexities add only finescale "noise" to a generally well behaved, though not unimodal, problem value considered as a function of the initial portfolio proportions (λ) to be optimised. This formulation as a low dimensional general nonconvex problem in the number of asset classes in the model is possible since for fixed λs subsequent rebalance decisions may be computed either by one iteration of nested Benders decomposition or directly by algebraic calculation (Dempster *et al.*, 2003).

6.3.2. In an attempt to reduce transaction costs it is also possible to define a model and corresponding trading strategy which rebalances to the fixed mix proportions only when current portfolio proportions have varied by more than specified percentages. Such a *relaxed fix mix* model has dead zones in which no portfolio rebalancing is necessary and may also be formulated as a global optimisation problem (using nested Benders decomposition) in the initial portfolio proportions.

6.3.3. To attack these nonconvex problems we have applied a variety of algorithms and software – local smooth approximate conjugate directions (Powell, 1964), the DIRECT global Lipschitz smooth partitioning algorithm (Gablonsky, 1998) and several others – to fixed mix variants of the CALM model with reasonable success (Scott, 2002). Currently we are working on improving the efficiency of these methods to make their routine operational use more robust.

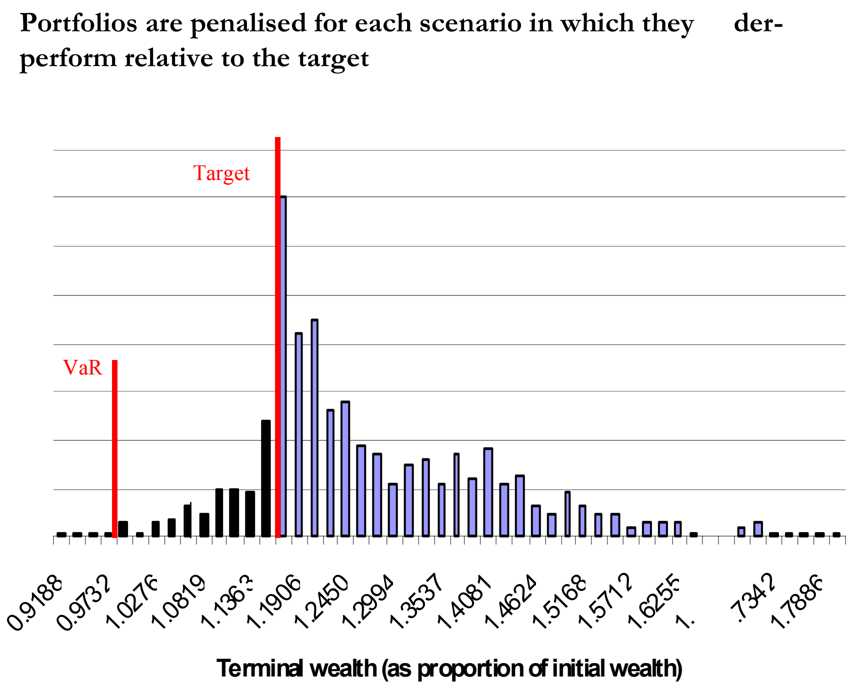

Figure 6.1 Terminal wealth distribution from a scenario tree

6.4. Capital Guaranteed Products Algorithm

6.4.1. The so-called *chance-constrained* programme arising from applying one or more probabilistic VaR-type capital guarantee constraints to the CALM model would only be convex if the distribution of current wealth $w_{t'}$ satisfies certain analytic conditions (Prékopa, 1980). This is not the case of course for a finite scenario-based distribution and hence the resulting problem is *nonconvex* and will require approximation for practical purposes. Like the benchmark portfolio problem however this approximation problem is not intractable. Instead of solving a problem (involving, for example, expected terminal wealth) of the form

$$\max E\mathbf{w}_{T+1}$$

subject to

$$\mathbf{x} \in \mathbf{X}$$
$$\mathbf{w}_{T+1} = g(\mathbf{x})$$
$$P\left(\mathbf{w}_{T+1} \geq w^*_{T+1}\right) \geq 1-\alpha,$$

we repeatedly solve

$$\max \mathrm{E}\left[\beta \mathbf{w}_{T+1} - (1-\beta)\left(\mathbf{w}_{T+1} - \tilde{w}_{T+1}\right)_-^2\right]$$

subject to

$$x \in X$$
$$\mathbf{w}_{T+1} = g(\mathbf{x}),$$

while searching for a value of target wealth $\tilde{w}_{T+1}$ for which the probabilistic constraint is satisfied. Alternatively, a severe downside linear penalty can be employed and this appears to be better at shaping wealth so as to reproduce scenario-based problem confidence levels out-of-sample using simulator or historical data, see Figure 6.1. We are currently perfecting this method for operational use with long horizon problems.

6.4.2. We have also tested a 0-1 mixed integer programming formulation in which the binary variables are used to count explicitly scenarios on which the guaranteed fund wealth is violated, but this approach currently appears intractable for anything but toy problems.

7. SYSTEM HISTORICAL BACKTESTS

7.1. *Implementation*

7.1.1. In a practical implementation of the dynamic stochastic optimization approach to strategic DFA a new problem is solved for each trading time, t=1,…,T, and the initial portfolios implemented. At each time t, the asset return and exchange rate model's parameters are re-estimated and re-calibrated using historical data up to and including time t, and the initial values of the simulated scenarios are given by the actual values of the variables at that time.

7.1.2. There are several reasons for implementing our approach in this manner. The first is that the actual value of the variables at t=2 are unlikely to coincide with any values of the variables in the simulated scenarios at t=2. If this is the case then the optimal investment policy will be undefined. The second and more important reason is that re-estimating and re-calibrating the simulator's parameters at each time t captures information in the history of the variables up to that point. Since the asset return and exchange rate model employed is only an approximation to the real dynamics, using the most recent history should improve the scenario simulation.

7.1.3. For a given problem formulation, the process of implementing the stochastic optimization approach at each trading time t can be represented by the following system diagram of Figure 7.1 (*cf.* Figure 1.1). Much of this system has been automated for the purposes of this research.

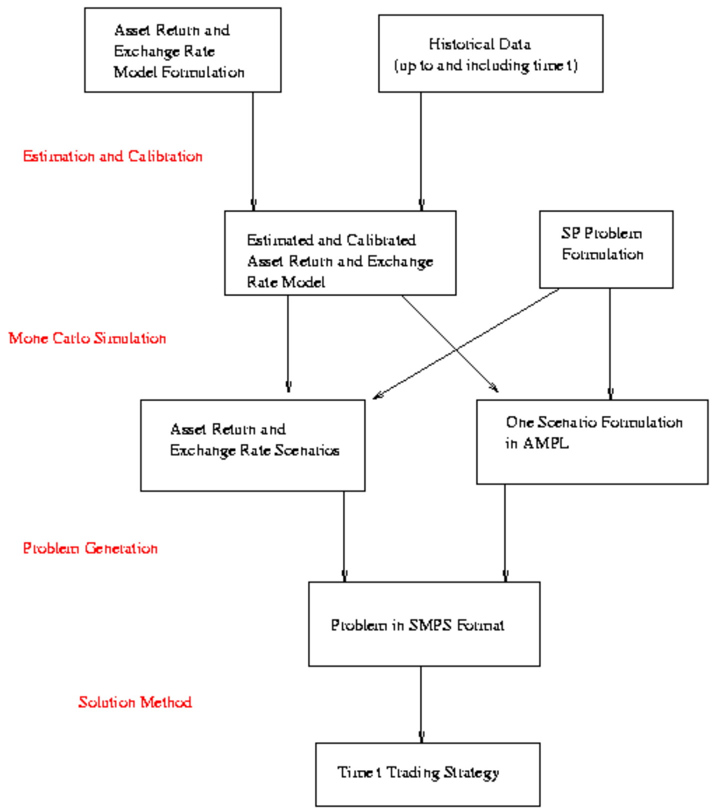

Figure 7.1 Pioneer CALM model system diagram

7.1.4. The quality of the first stage implementable and subsequent 'what-if' portfolio rebalance decisions of a dynamic stochastic optimisation ALM model clearly depend on a number of real world scenario contingent factors.

7.1.5. Obviously the most crucial factor is the 'predictive' power of the asset return statistical model underlying the scenario (tree) simulations. (We shall return to *ad hoc* tests of this factor in the next section). Less obvious perhaps is the impact of the *number* of scenarios used in the optimisation model and even more importantly the *branching structure* used in the scenario tree. Although there is general consensus that in dynamic models branching should be larger for the earliest decisions – in particular for the first implementable one – than for those later in the tree (see e. g. Dempster & Thompson, 2002) the number of scenarios required to stabilise problem value and decisions is highly model dependent. This is clearly a *sampling* problem for a continuous state stochastic optimisation problem – one level higher than a (discrete time) stochastic process sampling problem. Although asymptotic consistency results for both value and decisions are available (see Dempster (1998) or Shapiro (2002) and the references therein) the proofs are mathematically very difficult and the results of limited practical use. It is however generally agreed for a given problem that its value is stabilised by smaller scenario trees (samples) than are required to stabilise its (even implementable) decisions. Moreover, suppressing sampling error by the techniques discussed in §4.3 has also generally been seen to be beneficial for decision stability (although to an extent not reported in any detail in the literature). In our experiments, tree sizes (i.e. numbers of scenarios) have been reduced by a factor 5 by these means with a slightly greater problem run time reduction (to several minutes on a top end PC) which is of great practical use in fund design – although much remains to be done. We define a practical decision stability criterion in the next section.

7.2. *Backtesting*

7.2.1. Backtesting strategic DFA systems out-of-sample can take two forms: experimental and historical. In the more familiar *historical backtest*, statistical models are fitted to data up to a trading time t, scenario trees are generated to some chosen horizon t+T+1, the optimal decisions implemented at t are evaluated against historical returns at t+1, and the whole procedure rolled forward for T trading times. *Experimental backtests* can repeat this procedure as many times as is necessary to suppress sampling error by treating independently generated out-of-sample flat scenarios to T+1 as pseudo-histories. Such tests are invaluable in exploring the stability properties of decisions in specific models and we have termed a given model *decision stable* in scenario tree size and structure experiments when the standard deviation of the sampling error in each implementable decision portfolio proportion has been reduced to 10% of its sample mean value by a suitable choice of scenario tree for the model. Typically, 10,000 flat scenarios are used for such experiments.

7.2.2. Either type of backtest can involve a *telescoping horizon* as depicted in Figure 7.2 or a *rolling horizon* as shown in Figure 7.3.

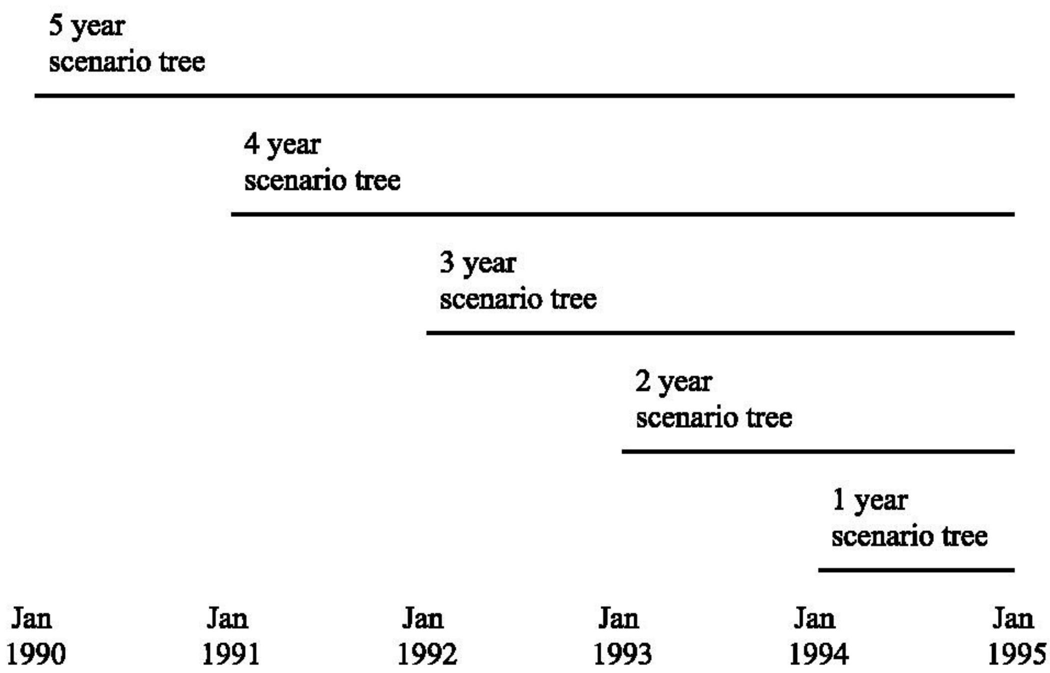

Figure 7.2 Telescoping horizon backtest schema

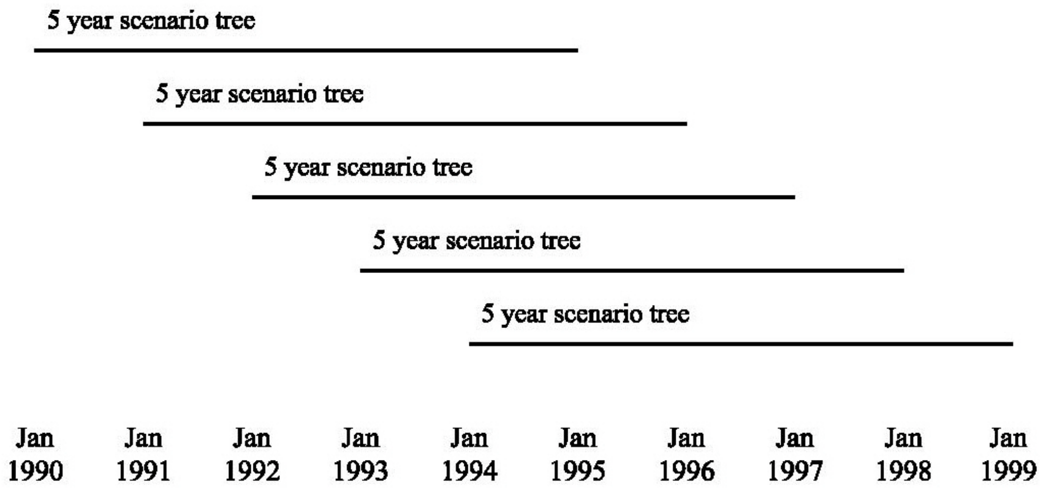

Figure 7.3 Rolling horizon backtest schema

7.3. Pioneer CALM Model Backtests

7.3.1. A number of historical backtests have been run on variants of the CALM global model, with perhaps surprisingly uniformly good results, see Villaverde (2002) for complete details. The aims of these tests were several. First we wished to establish the relative 'predictability' or otherwise of the alternative Pioneer tuned econometric models for short horizons. (Long horizon (20 or 30 year) experiments – where the simpler aim is merely to recapture historical statistical patterns – are currently in progress but require significant computational resources.) Secondly, we wished to understand the impact of the alternative utility functions available to the system on optimal portfolio decisions. Thirdly, we wished to evaluate the impact of risk attitudes imposed on fund wealth trajectories period-by-period (in terms of additively separable utility functionals) versus their imposition only on fund terminal wealth. Fourthly, we were interested in the farsightedness or otherwise of the dynamic stochastic optimisation approach to strategic DFA relative to rolling over single period-based systems à la Markowitz – the raison d'être of dynamic models. Finally, we were interested in what effects imposing the practical diversification and liquidity (turnover) constraints (T3 in Table 4.3) would have on backtest returns. We discuss the (at least partial) evidence to date on all these topics here.

7.3.2. All historical asset allocation backtests we report were from the viewpoint of a US dollar-based fund in Eire. The benchmark used is therefore the S&P500 equity index over the out-of-sample period for each test. All portfolio rebalances are subject to a 1% value tax on transactions which of course does not apply to the benchmark index. Monthly data (as set out in Table 3.1) were available from July 1977 to August 2002.

7.3.3. Figure 7.4 shows the results in terms of annualised returns of a typical backtest with a 2 year telescoping horizon and semi-annual rebalancing from February 1999 to February 2001 using the model of Appendix C with 8192 scenarios, a 128.16.2.2 branching structure and a terminal wealth criterion. During this period the S&P500 returned 0 percent. With no position limits the model tends to pick the best asset(s) and so in this case a high annual historical return is an indication of predictability in the tuned econometric model used to generate the scenarios. Once more realistic constraints are imposed in this test however portfolios become well diversified and in the results corresponding to the various attitudes to risk there is little to choose from. However, performance is improved by the use of the emerging market asset returns even though they were actually not used in the optimal portfolios. Corresponding results for the addition of the US economic model to the system are mixed. When this backtest was extended one period to August 2001 – when the S&P500 annualised return over the 2.5 year period was –2.3% – similar results were obtained with the best position limited result being 6.8% per annum for the downside-quadratic utility with a: = 0.5 and target wealth a 61% increase over the period.

	Capital Markets		Capital Markets + Emerging Markets		Capital Markets + Emerging Markets + US Economic Model	
Utility Function	No Limits	20% Limits	No Limits	20% Limits	No Limits	20% Limits
Linear	91%	9%	92%	10%	31%	11%
Quadratic	8%	9%	6%	11%	21%	6%
Downside-quadratic	54%	9%	70%	11%	29%	9%
Exponential	72%	9%	92%	10%	51%	11%
Power	91%		92%		49%	

Figure 7.4 Asset allocation backtests: Annualised returns from February 1999 – February 2001

7.3.4. Overall, the best overall historical backtest results were obtained using the downside-quadratic utility function with appropriate parameters. A summary of the backtests performed to date for this attitude to risk is given in Table 7.1. Note here that imposing the practical liquidity (T3) constraints, which could be expected generally to reduce returns, sometimes led to significantly increased returns. Notice also that the imposition of an attitude to risk of wealth in *each* period – the 10 year 5 year horizon rolling 4 area backtest using the linearised VARSIM simulation – improved annual return over the position limited returns for the two constituent 5 year periods (using 3 and 4 area capital market models) employing only an attitude to risk on fund terminal wealth.

7.3.5. Table 7.2 shows analysed implemented solver output for an historical backtest over the period 1996-2001 with annual rebalancing and the liquidity (T3) constraints imposed (corresponding to the bolded entry in Table 7.1). Note that the successive implemented portfolios are responding as much as possible to changing market conditions by asset allocations with varying diversification.

7.3.6. Overall, we found that the imposition of the T3 liquidity constraints in the model forced its decisions to take full advantage of the information in future scenarios and optimal forward rebalances to result in well diversified portfolios and significant improvement in historical backtest performance over rolling myopic single period models (*cf.* Hicks-Pedrón, 1998).

8. CONCLUSIONS

8.1. This paper describes an innovative joint project to construct the model base for a decision support system for defined contribution pension fund design at the strategic level. Each block of the system diagram of Figure 1.2 has been described in detail (including the third party component software utilised). The methods developed are much more widely applicable to a range of strategic DFA problems in finance. Practical solutions to two new problems – optimal fund benchmark setting and value-at-risk constrained guaranteed return fund design – have been outlined. In all historical backtests using data over roughly the past decade the global asset allocation system equalled or outperformed the S&P500 when transactions costs are taken into account. All system returns for the nonlinear statistical model were positive – even through the recent high tech crash.

8.2. A number of areas for further work have been identified throughout the paper and much work remains to be done. However if we have convinced the reader that the dynamic stochastic optimisation approach to strategic DFA problems is a practical reality today, the paper will have achieved its aims.

8.3. Currently we are developing an industrial strength version of the expected value of perfect information importance sampling algorithm (Dempster, 1998) represented by the inner dotted feedback loop in Figure 1.2. Eventually it should be possible to automate the reestimation and updating procedure of the outer dotted loop in the figure, but this adaptive filtering approach for this application is still a long way off.

8.4. The fund manager intends to become a leader in the management of pension funds for third parties. Its collaboration with the Centre for Financial Research at Cambridge has already made possible important advances in both its long-term forecasting engines and its optimisation techniques. Such know-how is currently used in the development of its new financial products and services.

9. ACKNOWLEGEMENTS

The authors would like to express their gratitude to the ALM Committee of Pioneer Investments and its parent company Unicredito Italiano for their continuing support. The project could not have been undertaken without the creativity and hard work of our colleagues (past and present): Santiago Arbeleche, Gabriele Cipriani, Giorgio Consigli, Maurizio Ferconi, James Scott and Giles Thompson, to whom we wish to record our heartfelt appreciation and thanks. We are grateful for the comments of the two anonymous referees which greatly improved the final version of this paper.

Initial Estimation Period	Out-of-sample Period	Length	Asset Return Model	Simulator	Number of Scenarios k	Rebalance Frequency	Risk Management Criterion	Horizon	Constraint Annualised Return % (see Section 5.4)			S&P 500 Benchmark Annualised Return %
									T1	T2	T3	
1972-1990	1990-1995	5 years	3 areas (ex Japan)	BMSIM	4	annual	terminal	telescoping	10.33	9.34	-	7.41
1992-1996	1996-2001	5 years	4 areas	BMSIM	4	annual	terminal	telescoping	13.36	7.13	-	14.12
1992-1996	1996-2001	5 years	4 areas	VARSIM	4	annual	terminal	telescoping	1.51	8.30	-	14.12
1992-1999	1999-2001	2.5 years	4 areas	BMSIM	8.2	semi-annual	terminal	telescoping	27.89	6.48	2.69	-2.30
1992-1999	1999-2001	2.5 years	above + emerging markets	BMSIM	8.2	semi-annual	terminal	telescoping	16.98	5.72	3.38	-2.30
1992-1999	1999-2001	2.5 years	above + US economy	BMSIM	8.2	semi-annual	terminal	telescoping	19.16	4.64	-0.38	-2.30
1992-1999	1999-2001	2.5 years	4 areas	VARSIM	8.2	semi-annual	terminal	telescoping	-6.40	-	-3.92	-2.30
1990-1996	1996-2001	5 years	4 areas	BMSIM	8.2	annual	all periods	telescoping	8.54	-	8.37	14.12
1990-1996	1996-2001	5 years	4 areas	VARSIM	8.2	annual	all periods	telescoping	5.78	9.99	**9.37**	14.12
1990-1996	1996-2001	5 years	4 areas	HSIM	8.2	annual	all periods	telescoping	4.95	-	6.04	14.12
1972-1991	1991-2001	10 years	4 areas	VARSIM	8.2	annual	all periods	5-year rolling	3.56	-	9.98	12.72

Table 7.1 Summary of CALM US$ fund historical backtests

VARSIM, 1996-2001, annual rebalancing, 8192 scenarios, additive downside-quadratic utility, T3 constraints															
	USstock	UScash	USbond	UKstock	UKcash	UKbond	UK fx	EUstock	EUcash	EUbond	EU fx	JPstock	JPcash	JPbond	JP fx
Date: Feb-96															
First Stage Weights	0.19	0.25	0	0.03	0	0.52	0	0	0	0	0	0	0	0	0
Historical return (dollar)	1.23	1.05	1	1.22	1.13	1.24	1.07	1.21	0.92	1.01	0.9	0.78	0.88	0.98	0.87
12-Month Portfolio Return Against History	1.18														
Date: Feb-97															
First Stage Weights	0.23	0.25	0	0	0	0.4	0	0	0	0	0	0.12	0	0	0
Historical return (dollar)	1.33	1.05	1.19	1.28	1.08	1.26	1.01	1.36	0.96	0.99	0.92	0.87	0.96	1.04	0.95
12-Month Portfolio Return Against History	1.17														
Date: Feb-98															
First Stage Weights	0.4	0.06	0.06	0	0	0.28	0	0.15	0	0	0	0.04	0	0	0
Historical return (dollar)	1.18	1.05	1.13	1.03	1.04	1.24	0.98	1.11	1.06	1.15	1.02	0.93	1.06	1.04	1.06
12-Month Portfolio Return Against History	1.16														
Date: Feb-99															
First Stage Weights	0.4	0	0.06	0	0	0.15	0	0.29	0	0	0	0.09	0	0	0
Historical return (dollar)	1.1	1.05	0.94	1.04	1.04	1	0.99	1.17	0.9	0.8	0.88	1.66	1.08	1.12	1.08
12-Month Portfolio Return Against History	1.15														
Date: Feb-00															
First Stage Weights	0.41	0	0	0	0	0	0	0.45	0	0	0	0.14	0	0	0
Historical return (dollar)	0.91	1.06	1.19	0.88	0.97	0.98	0.92	0.86	1.01	1.06	0.96	0.68	0.94	1	0.94
12-Month Portfolio Return Against History	0.85														

Table 7.2 Implemented annual portfolio rebalances for an historical backtest with liquidity constraints using VARSIM

REFERENCES

1. ARBELECHE, S. A. (2002). Econometric modelling for global asset liability management. Working Paper, Centre for Financial Research, University of Cambridge, in preparation.

2. BIRGE, J. R., DEMPSTER, M. A. H., GASSMANN, H. I., GUNN, E. A., KING, A. J. & WALLACE, S. W. (1986). A standard input format for multiperiod stochastic linear programs. *Committee on Algorithms Newsletter*, **17**, 1-20. Mathematical Programming Society.

3. BOENDER, G. C. E., van AALST, P. & HEEMSKERK, F. (1998). Modelling and management of assets and liabilities of pension plans in the Netherlands. In Ziemba & Mulvey (1998), *op. cit.*, 561-580.

4. BRADLEY, S. P. & CRANE, D. B. (1972). A dynamic model for bond portfolio management. *Management Science*, **19**, 131-151.

5. CAIRNS, A. J. G. (2000). A multifactor model for the term structure and inflation for long-term risk management with an extension to the equities market. Working Paper, Department of Actuarial Mathematics and Statistics, Heriot-Watt University.

6. CAMPBELL, R. (2000). The economic factor model: Theory and applications. Lehman Brothers presentation to Europlus Research and Management, Dublin, 31 March 2000.

7. CARIÑO, D. R., KENT, T., MYERS, D.H., STACY, C., SYLVANUS, M., TURNER, A. L., WATANEBE, K. & ZIEMBA, W. T. (1994). The Russell-Yasuda Kasai model: An asset/liability model for a Japanese insurance company using multistage stochastic programming. *Interfaces*, **24**, 29-49.

8. CHELLATHURAI, T. & DRAVIAN, T. (2002). Generalized Markowitz mean-variance principles for multi-period portfolio selection problems. *Journal of the Royal Society*, to appear.

9. COCHRANE, J. H. (1997). *Time series for macroeconomics and finance.* Unpublished book manuscript.

10. CONSIGLI, G. & DEMPSTER, M. A. H. (1998). The CALM stochastic programming model for dynamic asset-liability management. In Ziemba & Mulvey (1998), *op. cit.,* 464-500.

11. DANTZIG, G. B. & INFANGER, G. (1993). Multi-stage stochastic linear programs for portfolio optimization. *Annals of Operations Research,* **45**, 59-76.

12. DEMPSTER, M. A. H., ed. (1980). *Stochastic programming*. Academic Press.

13. DEMPSTER, M. A. H. (1988). On stochastic programming II: Dynamic problems under risk. *Stochastics,* **25**, 15-42.

14. DEMPSTER, M. A. H. (1993). CALM: A stochastic MIP model. Technical Report, Department of Mathematics, University of Essex.

15. DEMPSTER, M. A. H. (1998). Sequential importance sampling algorithms for dynamic stochastic programming. Working Paper 32/98, Judge Institute of Management, University of Cambridge. *Mathematical Programming*, to appear.

16. DEMPSTER, M. A. H., HICKS PEDRON, N., MEDOVA, E. A., SCOTT, J. E. & SEMBOS, A. (2000). Planning logistics operations in the oil industry. *Journal of the Operational Research Society*, **51**, 1271-1288.

17. DEMPSTER, M. A. H. & IRELAND, A. M. (1988). A financial expert decision support system. In *Mathematical models for decision support*. NATO ASI Series F48, G. Mitra, ed. Springer-Verlag, 415-440.

18. DEMPSTER, M. A. H., SCOTT, J. E. & THOMPSON, G. W. P. (2002). Stochastic modelling and optimization using STOCHASTICS™. In *Applications of Stochastic Programming.* S. W. Wallace & W. T. Ziemba, eds. MPS-SIAM Series on Optimization. Society for Industrial and Applied Mathematics, forthcoming.

19. DEMPSTER, M. A. H., SCOTT, J. E. & THOMPSON, G. W. P. (2003). Two nonconvex problems in strategic financial planning. Working Paper, Centre for Financial Research, University of Cambridge, in preparation.

20. DEMPSTER, M. A. H. & THOMPSON, G. W. P. (2002). Dynamic portfolio replication using stochastic programming. In *Risk management: Value at risk and beyond.* M. A. H. Dempster, ed. Cambridge University Press, 100-128.

21. DEMPSTER, M. A. H. & THORLACIUS, A. E. (1998). Stochastic simulation of international economic variables and asset returns: The Falcon Asset Model. *Proceedings of the 8^{th} International AFIR Colloquium*. Institute of Actuaries, London, 29-45.

22. DERT, C. L. (1995). A multi-stage stochastic programming approach to asset/liability management. PhD Thesis, Econometric Institute, Erasmus University.

23. DOAN, T. A. (1996). *RATS users manual*, Version 4.0.

24. DUVAL, D. B., TEEGER, M. H. & YAKOUBOV (1999). The TY model, a stochastic investment model for asset liability management. Technical Report, Aon Consulting, London.

25. FOURER, R., GAY, D. M. & KERNIHAN, B. W. (1993). *AMPL: A modelling language for mathematical programming*. Scientific Press.

26. GABLONSKY, J. (1998). An implementation of the DIRECT algorithm. Research Report, Centre for Research in Scientific Computation, North Carolina State University.

27. GARRATT, A., LEE, R., PESARAN, M. H., SHIN, Y. (2000). A structural cointegrating VAR approach to macroeconometric modelling. In *Econometric modelling: Techniques and applications.* S. Holly & M. Weale, eds. Cambridge University Press, Chapter 5.

28. GASSMAN, H. I. (1990). MSLiP: A computer code for the multistage stochastic linear programming problem. *Mathematical Programming*, **47**, 407-423.

29. GEYER, A. HEROLD, W. KONTRINER, K. & ZIEMBA, W.T. (2001). The Innovest Austrian pension fund financial planning model InnoALM. Technical Report, Siemens AG, Vienna.

30. GILL, P. E., MURRAY, W. & SAUNDERS, M. A. (2002). SNOPT: An SQP algorithm for large-scale constrained optimisation. *SIAM Journal of Optimization*, **12**, 979-1006.

31. HAKANSSON, N. H. (1974) Convergence to isolastic utility and policy in multiperiod portfolio choice. *Journal of Financial Economics*, **1,** 201-224.

32. HAMILTON, J. D. (1994). *Time series analysis*. Princeton University Press.

33. HICKS-PEDRÓN, N. (1998). Model-based asset management: A comparative study. PhD Thesis, Centre for Financial Research, University of Cambridge.

34. HODRICK, R. & VASSALOU, M. (2002). Do we need multi-country models to explain exchange rate and interest rate and bond return dynamics. *Journal of Economic Dynamics and Control*, **26**, 1275-1299.

35. HORNIMAN, M. D., JOBST, N. J., LUCAS, C. A. & MITRA, G. (2000). Constructing efficient portfolios: Alternative models and discrete constraints – a computational study. Technical Report, Department of Mathematical Sciences, Brunel University.

36. HØYLAND, K, KAUT, M. & WALLACE (2001). A heuristic for generating scenario trees for multistage decision problems. Working Paper, Molde University College, Molde, Norway.

37. HØYLAND, K. & WALLACE, S. W. (2001). Generating scenario trees for multistage decision problems. *Management Science,* **47**, 295-307.

38. IBM (1972). *Mathematical programming subsystem – Extended (MPSX) and Mixed Integer Programming (MIP) program description*. Document GH19-1091-0.

39. ILOG (2000). OPL Manual.

40. JARVIS, S., SOUTHALL, F. & VARNELL, E. C. (2001). Modern valuation techniques. Presented to the Staple Inn Actuarial Society, London, 6 February 2001.

41. KALLBERG, J. G. & ZIEMBA, W. T. (1983). Comparison of alternative utility functions in portfolio selection problems. *Management Science*, **29**, 1257-1276.

42. KAUFMANN, R., GADMER, A. & KLETT, R. (2001). Introduction to dynamic financial analysis. *ASTIN Bulletin,* **31,** 214-249.

43. KUSY, M. I. & ZIEMBA (1986). A bank asset and liability management model. *Operations Research,* **34**, 356-376.

44. KYRIACOU, M. N. (2001). Financial risk measurement and extreme value theory. PhD Thesis, Centre for Financial Research, University of Cambridge.

45. LANE, M. & HUTCHINSON, P. (1980). A model for managing a certificate of deposit portfolio under uncertainty. In Dempster (1980), *op. cit.*, 473-493.

46. MARKOWITZ, H. M. (1952). Portfolio selection. *Journal of Finance*, **7**, 77-91.

47. MEESE, R. & ROGOFF, K. (1983a). Empirical exchange rate models of the seventies. *Journal of International Economics*, **14,** 3-24.

48. MEESE, R. & ROGOFF, K. (1983b). The out-of-sample failure of empirical exchange rate models. In *Exchange rates and international macroeconomics*. J. Frenkel, ed. University of Chicago Press.

49. MERTON, R. C. (1990). *Continuous-time finance*. Blackwell Publishers.

50. MULVEY, J. M. (1989). A surplus optimization perspective. *Investment Management Review*, **3**, 31-39.

51. MULVEY, J. M. (1995). Financial planning via multi-stage stochastic optimization. Report SOR-94-09, Program in Statistics and Operations Research, Princeton University.

52. MULVEY, J. M. & THORLACIUS, A. E. (1998). The Towers Perrin global capital market scenario generation system. In Ziemba & Mulvey (1998), *op. cit.*, 286-314.

53. MULVEY, J. M. & VLADIMIROU, H. (1992). Stochastic network programming for financial planning problems. *Management Science,* **38**, 1642-1664.

54. PESARAN, M. H. & SCHUERMANN, T. (2001). Modeling regional interdependencies using a global error-correcting macroeconomic model. Working Paper, Trinity College, Cambridge.

55. POWELL, M. J. D. (1964). An efficient method for finding the minimum of a function of several variables without calculating derivatives. *Computer Journal*, **7**, 303-307.

56. PRÉKOPA, A. (1980). Logarithmic concave measures and related topics. In Dempster (1980), *op. cit.*, 63-82.

57. SCHERER, A. (2002). *Portfolio construction and risk budgeting.* Risk Books.

58. SCOTT, J. E. (2002). Modelling and solution of large scale stochastic programmes. PhD thesis, Centre for Financial Research, University of Cambridge.

59. SHAPIRO, A. (2002). Dynamic sampling schemes for multistage stochastic optimization. Research Report, Department of Operations Research, Georgia Institute of Technology.

60. SMITH, A. D. & SPEED, C. (1998) Gauge transforms in stochastic investment modelling. In *Proceedings of the 8th International AFIR Colloquium*. Institute of Actuaries, London.

61. SMITH, A. D. (1996). How actuaries can use financial economics. *British Actuarial Journal,* **2**, 1057-1193.

62. STEINBACH, M. C. (1999). Markowitz revisited: Single period and multi-period mean-variance models. Technical Report SC 99-30, Konrad-Zuse-Zentrum für Informationstechnik Berlin.

63. VANDERBEI, R. (2002). *Linear programming: Foundations and extensions*. Princeton University Press.

64. VILLAVERDE, M. (2002). Global fund management using stochastic optimization. Working Paper, Centre for Financial Research, University of Cambridge, in preparation.

65. WHITE, H. (1982). Maximum likelihood estimation of misspecified models. *Econometrica*, **50**, 1-25.

66. WILKIE, A. D. (1986). A stochastic investment model for actuarial use. *Transactions of the Faculty of Actuaries*, **39**, 391-403.

67. WILKIE, A. D. (1995). More on a stochastic asset model for actuarial use. *British Actuarial Journal*, **1**, 777-964.

68. WILKIE, A. D. (2000). The long term outlook for asset returns. Presented to the Inquire Europe Conference, Venice, 23 October, 2000.

69. ZENIOS, S. A. (1998). Asset and liability management under uncertainty for fixed income liabilities. In Ziemba & Mulvey (1998), *op. cit.,* 537-560.

70. ZIEMBA, W. T. & MULVEY, J. M. (1998), eds. *Worldwide asset and liability modelling.* Cambridge University Press.

APPENDIX A DETAILED PIONEER ASSET RETURN MODELS

Let the home currency be USD, and for a given scenario let:

- S_t^c denote the equity price at time t for c=US,UK,EU,JP,EM. S_t^{EM} is assumed to be denominated in USD
- R_t^c denote the percent return on lending cash between time t-1 and time t for c=US,UK,EU,JP. The percent return on borrowing cash is taken to be $R_t^c + \delta$ for some $\delta > 0$
- L_t^c denote the bond yield, expressed as a monthly percent, at time t for c=US,UK,EU,JP. The maturity and compounding frequency depends on c and is specified in Table 3.1
- X_t^c denote the exchange rate at time t for c=UK,EU,JP expressed as \$/local currency of c
- B_t^{EM} denote the EM bond price at time t denominated in USD; the (average) maturity of the bond (index) is specified in Table 3.1
- C_t^{US} denote the US Consumer Price Index (CPI) at time t
- W_t^{US} denote US wages at time t
- G_t^{US} denote US GDP at time t
- P_t^{US} denote US public sector borrowing (PSB) at time t.

The formulation of the capital markets discrete time model corresponding to the BMSIM3 simulator (see §4.4) is given by the following (with a monthly time step):

$$\frac{\boldsymbol{S_{t+1}^{US}} - S_t^{US}}{S_t^{US}} = \begin{pmatrix} a_S^{US} + a_{SS}^{US} S_t^{US} + a_{SR}^{US} R_t^{US} + a_{SL}^{US} L_t^{US} + a_{SX}^{US} X_t^{US} + \\ b_{SS}^{US} S_{t-1}^{US} + b_{SR}^{US} R_{t-1}^{US} + b_{SL}^{US} L_{t-1}^{US} + b_{SX}^{US} X_{t-1}^{US} \end{pmatrix} + \sigma_S^{US} \boldsymbol{\varepsilon_{St}^{US}}$$

$$\frac{\boldsymbol{R_{t+1}^{US}} - R_t^{US}}{R_t^{US}} = \begin{pmatrix} a_R^{US} + a_{RS}^{US}\left(\frac{S_t^{US}}{R_t^{US}}\right) + a_{RR}^{US}\left(\frac{1}{R_t^{US}}\right) + a_{RL}^{US}\left(\frac{L_t^{US}}{R_t^{US}}\right) + a_{RX}^{US}\left(\frac{X_t^{US}}{R_t^{US}}\right) + \\ b_{RS}^{US}\left(\frac{S_{t-1}^{US}}{R_t^{US}}\right) + b_{RR}^{US}\left(\frac{1}{R_t^{US}}\right) + b_{RL}^{US}\left(\frac{L_{t-1}^{US}}{R_t^{US}}\right) + b_{RX}^{US}\left(\frac{X_{t-1}^{US}}{R_t^{US}}\right) \end{pmatrix} + \sigma_R^{US} \boldsymbol{\varepsilon_{Rt}^{US}}$$

$$\frac{\boldsymbol{L_{t+1}^{US}} - L_t^{US}}{L_t^{US}} = \begin{pmatrix} a_L^{US} + a_{LS}^{US} S_t^{US} + a_{LR}^{US} R_t^{US} + a_{LL}^{US} L_t^{US} + a_{LX}^{US} X_t^{US} + \\ b_{LS}^{US} S_{t-1}^{US} + b_{LR}^{US} R_{t-1}^{US} + b_{LL}^{US} L_{t-1}^{US} + b_{LX}^{US} X_{t-1}^{US} \end{pmatrix} + \sigma_L^{US} \boldsymbol{\varepsilon_{Lt}^{US}}$$

$$\frac{\boldsymbol{S_{t+1}^{c}} - S_t^{c}}{S_t^{c}} = \begin{pmatrix} a_S^{c} + a_{SS}^{c} S_t^{c} + a_{SR}^{c} R_t^{c} + a_{SL}^{c} L_t^{c} + a_{SX}^{c} X_t^{c} + \\ b_{SS}^{c} S_{t-1}^{c} + b_{SR}^{c} R_{t-1}^{c} + b_{SL}^{c} L_{t-1}^{c} + b_{SX}^{c} X_{t-1}^{c} \end{pmatrix} + \sigma_S^{c} \boldsymbol{\varepsilon_{St}^{c}}$$

$$\frac{\boldsymbol{R}_{t+1}^c - R_t^c}{R_t^c} = \begin{pmatrix} a_R^c + a_{RS}^c\left(\frac{S_t^c}{R_t^c}\right) + a_{RR}^c\left(\frac{1}{R_t^c}\right) + a_{RL}^c\left(\frac{L_t^c}{R_t^c}\right) + a_{RX}^c\left(\frac{X_t^c}{R_t^c}\right) + \\ b_{RS}^c\left(\frac{S_{t-1}^c}{R_t^c}\right) + b_{RR}^c\left(\frac{1}{R_t^c}\right) + b_{RL}^c\left(\frac{L_{t-1}^c}{R_t^c}\right) + b_{RX}^c\left(\frac{X_{t-1}^c}{R_t^c}\right) \end{pmatrix} + \sigma_R^c \boldsymbol{\varepsilon}_{Rt}^c$$

$$\frac{\boldsymbol{L}_{t+1}^c - L_t^c}{L_t^c} = \begin{pmatrix} a_L^c + a_{LS}^c S_t^c + a_{LR}^c R_t^c + a_{LL}^c L_t^c + a_{LX}^c X_t^c + \\ b_{LS}^c S_{t-1}^c + b_{LR}^c R_{t-1}^c + b_{LL}^c L_{t-1}^c + b_{LX}^c X_{t-1}^c \end{pmatrix} + \sigma_L^c \boldsymbol{\varepsilon}_{Lt}^c$$

$$\frac{\boldsymbol{X}_{t+1}^c - X_t^c}{X_t^c} = \begin{pmatrix} a_X^c + a_{XS}^c\left(\frac{S_t^c}{X_t^c}\right) + a_{XR}^c\left(\frac{R_t^{US} - R_t^c}{X_t^c}\right) + a_{XL}^j\left(\frac{L_t^{US} - L_t^c}{X_t^c}\right) + a_{XX}^c\left(\frac{1}{X_t^c}\right) + \\ b_{XS}^c\left(\frac{S_{t-1}^c}{X_t^c}\right) + b_{XR}^c\left(\frac{R_{t-1}^{US} - R_{t-1}^c}{X_t^c}\right) + b_{XL}^c\left(\frac{L_{t-1}^{US} - L_{t-1}^c}{X_t^c}\right) + b_{XX}^c\left(\frac{1}{X_t^c}\right) \end{pmatrix} + \sigma_X^c \boldsymbol{\varepsilon}_{Xt}^c.$$

for c=UK,EU and X_t^{US}: $=X_t^{UK}$. The $\boldsymbol{\varepsilon}$ terms are correlated standard normal or standardized student t random variables. The a, b and σ terms are parameters of the model. Note that since we are assuming the home currency is USD, modelling an exchange rate for the US in unnecessary. Salient features of the model include non-linear drifts, a lag structure and constant volatilities in this form.

The formulation of the model for the BMSIM4 simulator is identical to that for BMSIM3 with the addition of Japan so that c=UK,EU,JP and with X_t^{US}: $=X_t^{JP}$. As noted in §3.2 additive binary dummy variables were used to remove the S_t^{JP} bubble and crash.

The formulation of the model for the BMSIM4EM simulator is identical to that for BMSIM4 with the addition of the following AR(1)/GARCH (1, 1) processes for EM equity and bonds

$$\begin{aligned}
&\frac{\boldsymbol{S}_{t+1}^{EM} - S_t^{EM}}{S_t^{EM}} = a_S^{EM} + a_{S1}^{EM}\frac{S_t^{EM} - S_{t-1}^{EM}}{S_{t-1}^{EM}} - a_{S2}^{EM}\sqrt{H_{t-1}^S}\varepsilon_{t-1}^S + \sqrt{H_t^S}\boldsymbol{\varepsilon}_t^S \\
&\boldsymbol{H}_t^S = b_S + b_{S1}H_{t-1}^S - b_{S2}H_{t-1}^S(\varepsilon_{t-1}^S)^2 \\
&\frac{\boldsymbol{B}_{t+1}^{EM} - B_t^{EM}}{B_t^{EM}} = a_B^{EM} + a_{B1}^{EM}\frac{B_t^{EM} - B_{t-1}^{EM}}{B_{t-1}^{EM}} - a_{B2}^{EM}\sqrt{H_{t-1}^B}\varepsilon_{t-1}^B + \sqrt{H_t^B}\boldsymbol{\varepsilon}_t^B \\
&H_t^B = b_B + b_{B1}H_{t-1}^B - b_{B2}H_{t-1}^B(\varepsilon_{t-1}^B)^2.
\end{aligned}$$

We assume that all $\boldsymbol{\varepsilon}$ terms are contemporaneously correlated but serially uncorrelated. Because the EM variables in BMSIM4EM only influence the US, UK, EU and JP financial variables via the shocks (the contemporaneously correlated ε terms) the EM variables will normally not influence the US, UK, EU and JP financial variables significantly.

The formulation for BMSIM4EME is similar to that for BMSIM4EM with the exception of two changes. The first is that we introduce the model for the US macroeconomic variables of §3.4. The second is that we the replace the US equations with:

$$\frac{\boldsymbol{S_{t+1}^{US}} - S_t^{US}}{S_t^{US}} = \begin{pmatrix} a_S^{US} + a_{SS}^{US} S_t^{US} + a_{SR}^{US} R_t^{US} + a_{SL}^{US} L_t^{US} + a_{SX}^{US} X_t^{US} + \\ b_{SS}^{US} S_{t-1}^{US} + b_{SR}^{US} R_{t-1}^{US} + b_{SL}^{US} L_{t-1}^{US} + b_{SX}^{US} X_{t-1}^{US} + \\ c_{SC}^{US} C_t^{US} + c_{SW}^{US} W_t^{US} + c_{SG}^{US} G_t^{US} + c_{SP}^{US} P_t^{US} + \\ d_{SC}^{US} C_{t-1}^{US} + d_{SW}^{US} W_{t-1}^{US} + d_{SG}^{US} G_{t-1}^{US} + d_{SP}^{US} P_{t-1}^{US} \end{pmatrix} + \sigma_S^{US} \boldsymbol{\varepsilon_{St}^{US}}$$

$$\frac{\boldsymbol{R_{t+1}^{US}} - R_t^{US}}{R_t^{US}} = \begin{pmatrix} a_R^{US} + a_{RS}^{US} \left(\frac{S_t^{US}}{R_t^{US}}\right) + a_{RR}^{US} \left(\frac{1}{R_t^{US}}\right) + a_{RL}^{US} \left(\frac{L_t^{US}}{R_t^{US}}\right) + a_{RX}^{US} \left(\frac{X_t^{US}}{R_t^{US}}\right) + \\ b_{RS}^{US} \left(\frac{S_{t-1}^{US}}{R_t^{US}}\right) + b_{RR}^{US} \left(\frac{1}{R_t^{US}}\right) + b_{RL}^{US} \left(\frac{L_{t-1}^{US}}{R_t^{US}}\right) + b_{RX}^{US} \left(\frac{X_{t-1}^{US}}{R_t^{US}}\right) + \\ c_{RC}^{US} C_t^{US} + c_{RW}^{US} W_t^{US} + c_{RG}^{US} G_t^{US} + c_{RP}^{US} P_t^{US} + \\ d_{RC}^{US} C_{t-1}^{US} + d_{RW}^{US} W_{t-1}^{US} + d_{RG}^{US} G_{t-1}^{US} + d_{RP}^{US} P_{t-1}^{US} \end{pmatrix} + \sigma_R^{US} \boldsymbol{\varepsilon_{Rt}^{US}}$$

$$\frac{\boldsymbol{L_{t+1}^{US}} - L_t^{US}}{L_t^{US}} = \begin{pmatrix} a_L^{US} + a_{LS}^{US} S_t^{US} + a_{LR}^{US} R_t^{US} + a_{LL}^{US} L_t^{US} + a_{LX}^{US} X_t^{US} + \\ b_{LS}^{US} S_{t-1}^{US} + b_{LR}^{US} R_{t-1}^{US} + b_{LL}^{US} L_{t-1}^{US} + b_{LX}^{US} X_{t-1}^{US} + \\ c_{LC}^{US} C_t^{US} + c_{LW}^{US} W_t^{US} + c_{LG}^{US} G_t^{US} + c_{LP}^{US} P_t^{US} + \\ d_{LC}^{US} C_{t-1}^{US} + d_{LW}^{US} W_{t-1}^{US} + d_{LG}^{US} G_{t-1}^{US} + d_{LP}^{US} P_{t-1}^{US} \end{pmatrix} + \sigma_L^{US} \boldsymbol{\varepsilon_{Lt}^{US}}.$$

The addition of the US macroeconomic variables is an attempt to create a more realistic model for asset returns and exchange rates. Because they influence the US financial variables through the drift terms and the shocks they should have a significant impact on the US financial variables. Again we assume that all $\boldsymbol{\varepsilon}$ terms are contemporaneously correlated but serially uncorrelated.

The generation of the dynamic stochastic optimisation problems requires the asset returns and exchange rates in each scenario. Appendix B explains how bond yields are transformed into bond asset returns.

APPENDIX B DERIVATION OF BOND RETURNS FROM BOND YIELDS

The following is a derivation of the 1 month bond return for the US. The UK, EU and Japan formulas differ only in the maturity and compounding frequency of the bond yield.

The US bond has a 30 year maturity with semi-annual compounding. Let $L1_t$ denote the 30 year annualised bond yield with semi-annual compounding, i.e. $L1_t = 12L_t/100$. Let F denote the face value of the bond, and let c_t denote the annual coupon rate.

Consider holding a newly issued 30 year bond from time t to time t+1 which is 1 month later. The value of the investment at time t is the cash price of the 30 year bond which is given by:

$$V_t = \sum_{n=1}^{60} \frac{F\frac{c_t}{2}}{(1+\frac{L1_t}{2})^n} + \frac{F}{(1+\frac{L1_t}{2})^{60}}$$

$$= \frac{Fc_t}{L1_t}(1-\frac{1}{(1+\frac{L1_t}{2})^{60}}) + \frac{F}{(1+\frac{L1_t}{2})^{60}} \quad .$$

At time t+1 or 1 month later there has been no coupon payment and the value of the investment is the cash price of a bond with a $29\frac{11}{12}$ year maturity and which pays a coupon in 5 months and then every 6 months until maturity. The cash price of this bond is:

$$\hat{V}_{t+1} = \sum_{n=1}^{60} \frac{F\frac{c_t}{2}}{(1+\frac{\widehat{L1}_{t+1}}{2})^{n-1+\frac{5}{6}}} + \frac{F}{(1+\frac{\widehat{L1}_{t+1}}{2})^{59\frac{5}{6}}} \quad .$$

If we assume that the yield of this bond $\widehat{L1}_{t+1} = L1_{t+1}$, we can approximate $\hat{V}_{t+1}$ by:

$$\tilde{V}_{t+1} = \sum_{n=1}^{60} \frac{F\frac{c_t}{2}}{(1+\frac{L1_{t+1}}{2})^{n-1+\frac{5}{6}}} + \frac{F}{(1+\frac{L1_{t+1}}{2})^{59\frac{5}{6}}}$$

$$= \frac{Fc_t}{L1_{t+1}(1-\frac{L1_{t+1}}{2})^{-\frac{1}{6}}}(1-\frac{1}{(1+\frac{L1_{t+1}}{2})^{60}}) + \frac{F}{(1+\frac{L1_{t+1}}{2})^{59\frac{5}{6}}} \quad .$$

Then the 1 month bond return can be estimated as:

$$\frac{\tilde{V}_{t+1}}{\tilde{V}_t} - 1 = \frac{\frac{c_t}{L1_{t+1}(1-\frac{L1_{t+1}}{2})^{-\frac{1}{6}}}(1-\frac{1}{(1+\frac{L1_{t+1}}{2})^{60}}) + \frac{1}{(1+\frac{L1_{t+1}}{2})^{59\frac{5}{6}}}}{\frac{c_t}{L1_t}(1-\frac{1}{(1+\frac{L1_t}{2})^{60}}) + \frac{1}{(1+\frac{L1_t}{2})^{60}}} - 1 \quad .$$

The coupon rate c_t can be approximated as some fraction of $L1_t$, i.e. $c_t = mL1_t$, with $m \le 1$.

APPENDIX C THE PIONEER CALM ASSET ALLOCATION MODEL

The mathematical formulation of the basic asset management problem in deterministic equivalent form for solution is given by the following version of the CALM model of Dempster (1993). We assume that u is given by one of the utility functions described in §5.2, that as a consequence of Monte Carlo simulation each scenario ω in Ω is equally likely, that there are no cash inflows or outflows and that the only regulatory and performance constraints are cash borrowing limits, short sale constraints, position limits and turnover constraints. Liabilities are easily added in terms of cash inflows or outflows (Consigli and Dempster, 1998).

$$\max_{\theta} \sum_{\omega\in\Omega} p(\omega)\sum_{t=2}^{T+1} u_t(w_t^{\theta}(\omega))$$

$s.t.$

$$w_1 + \sum_{i\in I} p_{i1}(\omega)(gx_{i1}^-(\omega) - fx_{i1}^+(\omega)) + \sum_{k\in K} p_{k1}(\omega)(-z_{k1}^+(\omega) + z_{k1}^-(\omega)) = 0 \qquad \omega\in\Omega$$

$$\sum_{i\in I} p_{it}(\omega)(gx_{it}^-(\omega) - fx_{it}^+(\omega)) +$$

$$\sum_{k\in K} p_{kt}(\omega)((1+r_{kt}^+(\omega))z_{kt-1}^+(\omega) - (1+r_{kt}^-(\omega))z_{kt-1}^-(\omega) - z_{kt}^+(\omega) + z_{kt}^-(\omega)) = 0 \qquad t = 2,...T, \omega\in\Omega$$

$$x_{i1}(\omega) = x_i + x_{i1}^+(\omega) - x_{i1}^-(\omega) \qquad i\in I, \omega\in\Omega$$

$$x_{it}(\omega) = x_{it-1}(\omega)(1+v_{it}(\omega)) + x_{it}^+(\omega) - x_{it}^-(\omega) \qquad i\in I, t = 2,...,T, \omega\in\Omega$$

$$\sum_{i\in I} p_{it}(\omega)(1+v_{it}(\omega))x_{it-1}(\omega) +$$

$$\sum_{k\in K} p_{kt}(\omega)((1+r_{kt}^+(\omega))z_{kt-1}^+(\omega) - (1+r_{kt}^-(\omega))z_{kt-1}^-(\omega)) = w_t^{\theta}(\omega) \qquad t = 2,...,T+1, \omega\in\Omega$$

$$W_t^{\theta}(\omega) = \sum_{j\in I} p_{jt}(\omega)x_{jt}(\omega) + \sum_{k\in K} p_{kt}(\omega)(z_{kt}^+(\omega) - z_{kt}^-(\omega)) \qquad t = 1,...,T, \omega\in\Omega$$

$$p_{kt}(\omega)z_{kt}^-(\omega) \le \bar{z}_k \qquad k\in K, t = 1,...,T, \omega\in\Omega$$

$$p_{it}(\omega)x_{it}(\omega) \ge \bar{x}_i \qquad i\in I, t = 1,...,T, \omega\in\Omega$$

$$p_{it}(\omega)x_{it}(\omega) \le \phi_i W_t^{\theta}(\omega) \qquad i\in I, t = 1,...,T, \omega\in\Omega$$

$$p_{kt}(\omega)(z_{kt}^+(\omega) - z_{kt}^-(\omega)) \le \phi_k W_t^{\theta}(\omega) \qquad k\in K, t = 1,...,T, \omega\in\Omega$$

$$| p_{it}(\omega)x_{it}(\omega) - p_{it-1}(\omega)x_{it-1}(\omega) | \le \alpha_i W_t^{\theta}(\omega) \qquad i\in I, t = 1,...,T, \omega\in\Omega$$

$$w_t^{\theta}(\omega) \ge 0 \qquad t = 2,...,T+1, \omega\in\Omega$$

$$x_{it}^+(\omega), x_{it}^-(\omega), z_{kt}^+(\omega), z_{kt}^-(\omega) \ge 0 \qquad i\in I, k\in K, t = 1,...,T, \omega\in\Omega,$$

where $p(\omega) = 1/|\Omega|$ and $\theta_{ikt}(\omega) := (x_{it}(\omega), x_{it}^{+}(\omega), x_{it}^{-}(\omega), z_{kt}^{+}(\omega), z_{kt}^{-}(\omega))$ for i in I, k in K, t=1,..., T, ω in Ω .

The first set of constraints are known as *cash balance* constraints. They insure that the net flow of cash at each time and in each state is zero. The next set of constraints are known as *inventory balance* constraint*s*. They give the position in each equity and bond asset at each time and in each state. The third set of constraints define respectively the before and after rebalancing wealth at each time in each state. The next six constraints are the cash borrowing constraints, short sale constraints, position limit constraints, turnover constraints and solvency constraints discussed in the previous section. This deterministic mathematical programming problem is convex, linearly constrained and (unless u is the identity) has a non-linear objective.

Global Asset Liability Management

ABSTRACT OF THE DISCUSSION

Mr M. Germano (introducing the paper)**:** Firstly, I should like to thank all of you for this exciting opportunity to present the results of the last three years of our work, in partnership with Professor Dempster and the financial research laboratory at Cambridge. I will give you a brief introduction to the background and objectives behind the system development, to highlight the true dilemma that we have been facing as an industry and as a fund management firm, and then discuss how we have been trying to cope with it.

In United Kingdom, the recent market, as you will know, has been characterised by extremely high volatility. We have seen nevertheless a growing appetite for long-term retirement plans in continental Europe. Clients have been attracted to products with either a minimum-return guarantee or a target-return with an associated level of probability of achievement. The general problem is familiar to all of you. That is, the problem of trying to fill in what we call the "Pensions Gap" by attempting to support the first and the second Pillars with products that will allow us to close the gap in order to maintain current standards of living for the new pensioners.

If you look at the problem from the client's point view, it is very difficult to identify the client with a single paradigm. Consequently, we spent a lot of time trying to understand their true needs. We will outline how it was difficult to associate all the clients with the standard "life cycle model", in which the employees keep contributing to have enough money to maintain their standard of living until death. We have a second model in which there is a true interest from the employee in contributing in order to build a final endowment to leave future generations. So a second model that we can identify is the "legacy model". There is a third model, the true entrepreneur, whose interests and risk aversion do not follow the usual path, but instead follow a path in which the risk aversion does not decrease with age and who retains an appetite for risky assets during the later stages of life.

So we have very complex needs. We try to answer with a complex model, that allows us to solve what we call the pensions plan dilemma. On one hand we keep trying to optimise long-term returns, and on the other side, minimise the downside risk over an intermediate horizon. Being able to cope with multiple schemes, in terms of contribution, is a necessity because we have some clients who may wish to make their contributions at the beginning through a lump sum investment, others, who wish to make regular contributions and each of these will have different risk aversion and different target for expected returns and minimum guaranteed returns.

These are the ingredients that we need to consider when offering complete customised products. Such solutions will give the opportunity to some further soft aspects, such as human capital. A young employee of 25 years of age is easily identifiable as a long-term investor. If we take into account factors like human capital, such as job security and the risk of being fired, it is clear that this potential future pensioner has different characteristics and may have some intermediate horizon, say over five years or six years, during which he wants to secure a minimum return at a certain level of probability in order to cope with such adverse events.

Professor M. A. H. Dempster, Hon.F.I.A. (introducing the paper)**:** I suppose that as leading the quantitative attack on the development of the system of this paper, I should advise that the paper is really a case study to investigate two aspects of general problems in asset management or asset liability management: uncertainty -- stochastics -- and even more importantly dynamics. In other types of dynamic financial modelling these two aspects appear. For example, in the banking industry at this time there is a great deal of interest in dynamic capital allocation against the risks of various business units. Again, the technology of this paper can be employed.

So, how does this technology differ from ones that actuaries are more familiar with? If actuaries try to optimise asset management over a fixed horizon, with or without liabilities, one typically use some kind of fairly static model, usually after generating -- perhaps very complex generation -- dynamic scenarios of capital market variables, of economic variables, demographic variables, and so on. Then they use some kind of discounting, deterministic or stochastic, to produce present values and finally employ static optimisation using essentially Markowitz technology. It might be surplus that is optimised in an insurance application, but it is still static.

Something I want to emphasise -- and the paper gives a great many details about this case, but this is the basic point -- is that it is better to use dynamic models. Nobody says that they are simple; but 30 years ago Markowitz models were not simple either, and now everyone uses Markowitz models. Since then computers have advanced. Mathematics have advanced. Statistics have advanced. Stimulation has advanced. I would argue that it is now time to employ these models because they offer very significant advantages.

Another way of looking at this is that in leading applications in the actuarial and financial worlds people currently make very complex simulations, which in the actuarial world, usually go by the name dynamic financial analysis. I referred to scenario generation earlier. These simulations, however, employ very complex models and have many parameters. They are extremely difficult to fit to the data available and it is extremely difficult to optimise their decision parameters price fit, because the typical method is simply to take a few parameters, step them through various values and then try to analyse the output, which is quite bewildering.

So, from that point view we have coined the term *strategic* dynamic financial analysis, meaning that we will add industrial strength optimisation to this recipe.

To deal with a system like this is extremely complex. In figure 1.1 I want to stress the left-hand side of the diagram, which shows that one has to put a lot of different skills together. The computer science which deals with data, that is raw statistical analysis of data, cleaning it, filtering it, and so on. Next comes building the econometric models for both asset returns (as begun in this country and, indeed, generally by Professor Wilkie) and possibly complex liability models, although we are not going to stress these in this paper. Then, once we have these models fitted and tuned to data, we must use stochastic Monte Carlo stimulation, which is dynamic financial analysis and take us essentially down to the end of the third bar. Below that is what we want to add to the process; namely, the handling of essentially large-scale dynamic financial analyses models with parameters and optimisation of those parameters. They might be portfolio balances. They might be certain decisions such as capital allocations or more complicated business decisions. One needs complex software for this, and there is quite a bit about this of the problem in the paper.

To go to building the scenario generators, we have in fact (and there are details in the paper of the models used, and so on) a global system -- hence the title of the paper -- that deals with the four major currency areas and emerging markets. We do this in a way that has been pioneered by Professor Wilkie and others; to build a canonical model for a currency area and then link them together. We have linked the areas together with foreign exchange on the grounds that rates react faster than, for example, trade flows, which would come to mind initially for a macro economist.

What makes our dynamic financial analysis approach different from usual dynamic financial analysis. In dynamic financial analysis this is a typical case of scenario generation involving the New York Stock Exchange Valuated Weighted (refer figure 4.6) Index. What we do here is set up our econometric model, or use some stochastic differential equation model with parameters which we tune until we feel that the model fits the data represented by the black line. Then we generate scenarios forward out of sample. These are the coloured paths in the figure. We look initially to see if they have the kind of variants, and so on, variability, volatility as the black path, which is the history, and they have a suitable distribution around historical path, and so on. All a bit of art rather than science, however, broadly used these days.

On the other hand, we do not just do this for one variable. In our model using 15 variables and throwing in the economic variables, we can have a system of up to 33 state variables. So this is already fairly large-scale. However, the key point about stochastic optimisation is that it is not the optimisation problem necessarily that is now so difficult, because there are industrial strength optimisers which used cleverly can double running speed or even raise it two orders of magnitude using decomposition techniques. Rather it is that one has to generate high dimensional scenarios in a conditional mode (figure 4.7). This is just a scheme that represents the dynamic decisions every node this tree represents, for example, a portfolio balance. Every such decision must face alternative real-world scenarios. That is what makes it different. Even the scenario generation is thus more complex than basic dynamic financial analysis and requires extra software although in principle is straightforward.

The second type of problem we have solved, and generated software to address, concerns the formulation and solution of asset liability management problems, particularly those of defined contribution pension plans. These products are essentially mutual funds with a guarantee or a target return that is achieved at least with high probability. We use utility functions here not in the sense of individuals, but to represent the risk attitude of perhaps a sort of representative individual for a specific fund, tailored with respect to a specific horizon and specific risk characteristics, to which Mr Germano alluded to.

However, for scenario based models, such problems with a probabilistic constraint pose a new scientific class, which we believe we have solved in the applications that we have looked so far. There is still a lot to do. Rather than use a simplified theoretical model, say, with independent returns, period to period, and other restrictive assumptions, we have been able to solve these problems when the scenario generator is of the complex type, which I showed you.

The third type of problem that we looked at was the idea of setting up benchmark portfolios optimally. I have just come back from Princeton and our colleagues there have worked on a similar type of problem as an asset allocation strategy. Here what are trying to do, since these problems define a restriction of the full optimum, and impose a lower bound on it, is use this bound and the associated returns as a benchmark against which other strategies, including the full optimum, can be measured.

What happens for general cases of all three of these classes of models is that one obtains other suitable formulations really a very large-scale deterministic problem. Essentially, every node of the tree, that I showed you, has assigned to it a set of asset allocations or capital allocations, depending on the application (here it is asset allocation) that correspond to the states, that is the values of the state variables, at the time. That can be expanded so that we have, for every such node a decision, what is represented in the deterministic model are the constraints which are shown and enumerated in section 2.23. These are usually linear, because accounting is usually a linear operation. However, the attitude to risk, which is embodied for a particular fund and a particular horizon, can be represented by either a sum of risk attitudes represented period by period by a utility function or in fact just looking at risk in the terminal wealth achieved at the horizon.

What we are doing is designing tools to shape the wealth distribution at a particular point in time as the fund evolves (refer figure 6.1). We are interested in not failing to meet a guarantee, for example, by calculating value at risk through accounting scenarios which fail to meet the guarantee. We will use here tens of thousands of scenarios, possibly, and always we can trade off a kind of return variable with the risk variable to shape the wealth distribution. We have put a lot of work into understanding this trade off and in effecting it technically.

We have also done a lot of backtesting of the systems and are very pleased with the backtests. They are as honest as they can be and the returns have been quite good relative to the S&P 500, which we used as a benchmark because we treated this as a dollar fund. These results may not really mean anything due to the small size, as I am sure many people will point out, but it is very nice to have good results rather than bad results, even through the peak and the crash of the world's markets.

The last point is that the kind of systems described require pretty fancy graphics, because one has to go back and forth between tweaking data, tweaking models, optimising, and so on. We will never get away from that process. As the models get more complicated, we may have to do a little more. Figure D1 shows a sort of prototype screen which we are redesigning and implementing in practice.

To conclude, strategic asset liability management in tactical risk systems which are dynamic are a reality today. Multiperiod models yield multiple advantages. They tend to have a more stable portfolio path as we roll them forward because they hedge against a great deal of alternative scenarios, both good and bad. We can get best, worst and value at risk limited portfolio views, or asset and liability views, looking down the scenarios so that we can do a lot of 'what-if' analyses with these kinds of models. The assumptions, that are needed to make rolling over of myopic models correct, are almost entirely violated by real market data.

Using the type of software systems, which are used now in industrial optimisers, both linear and non-linear, we can model any kind of constraint structure, including regulatory and tax structures, and so on. The dynamic stochastic, strategic DFA, approach results in very large models, because they are contingency plans. All possibilities are being considered simultaneously, so they involve literally millions of equations and variables, but they can be solved in a few minutes on a current PC.

Flexibility and visualisation are keys to providing effective decision support systems of this type for strategic pension planning, and the other financial applications that I mentioned (see Figure D1).

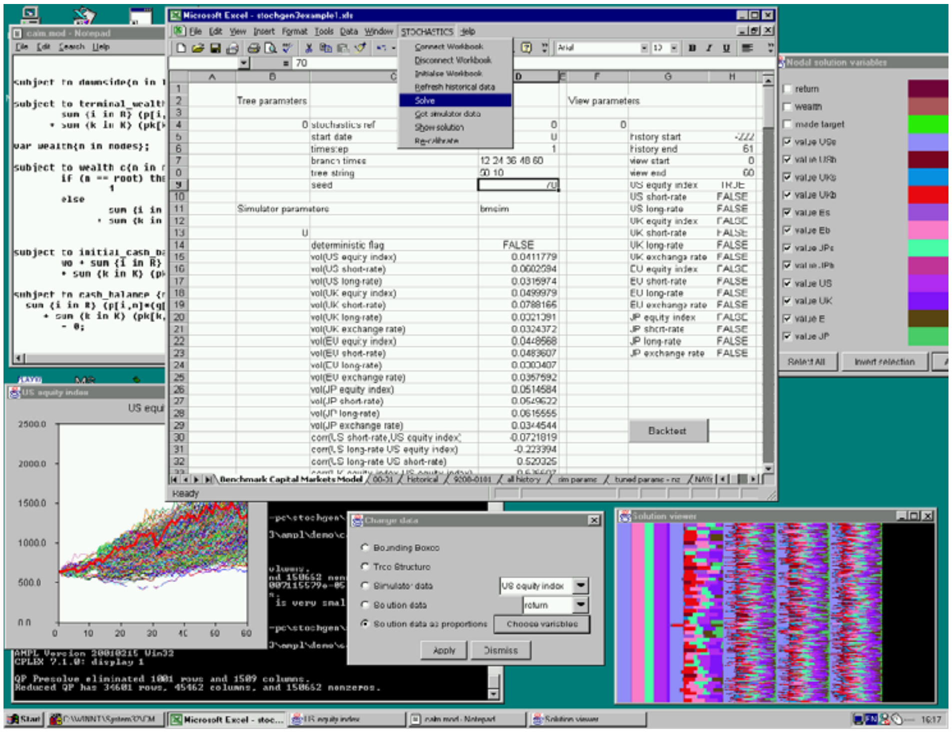

Figure D1 Prototype User Interface for the Fund Manager STOCHASTICS™ System Stochgen 3.0

Dr A. J. G. Cairns, F.F.A. (opening the discussion): Dempster, Germano, Medova and Villaverde's paper is very interesting and comprehensive with many stimulating ideas contained within it, which the wider profession needs to be learning about.

The paper assumes a certain knowledge of the subject of dynamic optimisation. It may be helpful to cover some of the basic ideas in this subject area. A very simple example will illustrate some of the basic ideas, which will provide a firmer base from where we can look at the more complex environment that is in the paper.

In this simple example we have two assets: cash and equities. Investment is carried out over two periods. We have to decide what proportion of our fund, or what I will from now on call wealth, should be invested in equities over each of the two investment periods. So p_0 represents the proportion in equities from time 0 to time 1 and p_1 the proportion in equities from times 1 to 2.

Our investor has a utility function U which he will apply to his wealth at time 2. His objective is to maximise his expected utility by choosing p_0 and p_1 optimally.

One way of going about this is to find the optimal static investment strategy, that is, we assume that p_0 and p_1 are equal and then optimise over p_0.

On the other hand, you have the opportunity at time 1 to revise your strategy in light of what has happened over the previous year.

What are the consequences of this? Well, first it makes sense to take advantage of this opportunity. If you do not, then you may be acting suboptimally, and you know you are acting potentially suboptimally by not taking this opportunity.

Second, how much you invest in equities at time 0 should take into account how you will act at time 1. This is the dynamic element.

Now compare this with the static investment strategy. The second approach must be better in terms of optimising in advance the expected utility because the first, static approach is just a special case of the second.

What we can also point out at this stage is that p_1 does not need to be specified at time 0. Instead the optimal equity proportion can and should take account of what has happened in the first year. So it might be random, but we will have a rule, or a set of rules, which will allow us to determine p_1 at the time when we need it.

This simple problem leads us to the Bellman Principle, which tells us that any optimisation problem over T years can be broken down into a sequence of T one-year optimisation problems. Each one-year problem (or whatever time-step you want) is effectively a static utility maximisation problem which is clearly much easier to tackle.

As a side remark, in continuous rather than discrete time we use the Hamilton-Jacobi-Bellman equation (or HJB equation). There are growing numbers of papers on this in the international actuarial journals such as *ASTIN Bulletin* and *Insurance: Mathematics and Economics*. Typically these papers are more concerned with the development of qualitative results using simpler models, than we have before us today. However, the analytical nature of their results do allow us to quantify exactly how much popular, but suboptimal, strategies cost to the policyholder in terms of expected utility.

Coming back now to the discrete-time problem, in some cases, for example, power utility, the static optimal solution is the same as the dynamic optimum. In other cases, the static optimum delivers a far inferior optimal expected utility to an investor than the dynamic optimum. It is in this sort of case that we need to be able to identify and solve the dynamic problem as accurately as possible.

The Bellman Principle requires what we call a finite-dimensional state variable, $X(t)$, that is, how we predict the future depends on a finite number of variables observable at the current time. For example,

$X_1(t)=W(t)$=current wealth

$X_2(t)=r(t)$=short-term interest rate

$X_3(t)=DY(t)$=equity dividend yield

$X_4(t)=I(t)$=inflation rate

etc.

In the simplest case we have only one state variable which is relevant and the investment problem is simple to optimise.

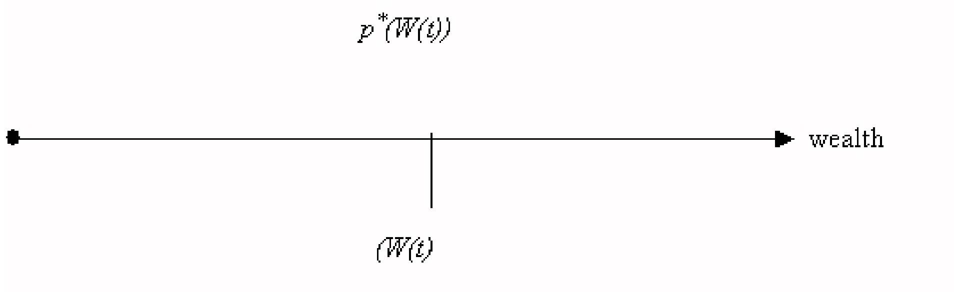

Figure D2

However, it is still not trivial, because we need to establish the optimal *p** for each possible wealth at time *t* as I have tried to show in Figure D2.

As the number of state variables increases things get much more complex as we need to find the optimal asset mix for each combination of the state variables as we see in Figure D3.

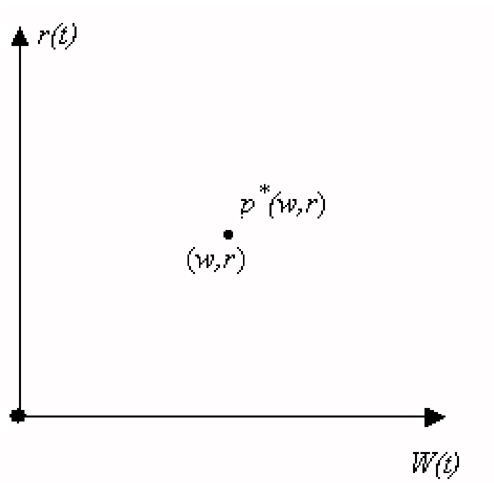

Figure D3

In today's paper we find ourselves in the situation where

(a) we want to carry out dynamic optimisation

(b) we have a large number of state variables.

Indeed, to add to the complexity of the problem the authors include transaction costs. This, then, also increases the number of state variables.

For example, think back to my very simple example with two assets.

What the investor will choose to do at time 1 will depend not just on

(a) total wealth at time 1 (as before), but also

(b) how this wealth is distributed between cash and equity at the end of the first period of investment.

The reason is that the effective wealth available depends on how much trading you want to do at time 1 and what the transaction costs of this are.

So we have a complex optimisation problem, particularly in terms of computer programing and processing times. The paper brings to us a wealth of expertise built up over many years in the field of dynamic stochastic optimisation. The authors have brought together and summarised a number of technically difficult toolkits to address this problem.

The other important part to this paper is the authors' presentation of their asset model. This gives us the state variables for the dynamic optimisation problems. However, the model is very interesting to see in its own right given the relative scarcity of other published models. It would have been nice to have typical parameter values for this model to allow us to get a feel for how it compares with other models.

Some specific observations now and questions, which the authors may wish to respond to during the meeting.

In ¶3.4.3 the authors describe a possible term-structure model and I have some thoughts and questions on this.

(1) What is the difference between R^0 and R?

They need to be different because of the use of the process Y and its dependence on a third dZ term.

(2) Both R^0 and L can go negative, so how do you deal with this in practical terms?

(3) It appears that the real world version of these processes follows a random walk with drift through the constants δ. Could the authors comment on this, because my interpretation seems rather unsatisfactory and therefore probably incorrect.

(4) Finally how are we to interpret the α's, the market prices of risk? For example, how do they translate into risk premiums on various bonds?

In ¶3.1 the authors describe the different types of asset model: in particular, those developed in the econometric-modelling school and those developed in the arbitrage-free modelling school. To me the paper polarises things in a slightly misleading way. That is, a given model will belong to one school or the other with no middle ground. However, I can take my own model, cited in this section, as an example. This model was indeed developed starting from an arbitrage-free modelling standpoint. The approach implied in today's paper is to calibrate the model using today's market prices with a view to getting the best possible estimate of the market price of, say, some insurance or pensions liabilities. On the other hand, most of the arbitrage-free models can be calibrated using historical data. In particular, we can use exactly the same statistical methods of estimation as employed in fitting the econometric models. Thus there is no *a priori* reason why econometric models should be better than arbitrage-free models for long-term risk management purposes. Indeed, if a model is being used for dynamic optimisation, then it must be arbitrage free otherwise dynamic optimisation, properly implemented, will find the arbitrage to create infinite profit. It might well be that the authors feel that typical, arbitrage-free models are unnecessarily complex within the context of a global asset model, and this is true, but this is quite a different issue from the point that is brought out in that particular section.

At a couple of points in the paper the authors describe the use of a branching structure in the scenario generator. In particular, in ¶7.3.3 they use a 128x16x2x2 structure. I was quite intrigued by this. It is obvious that having two branches is quite inadequate for determining the optimal strategy in the final step, so the natural conclusion, perhaps, is that we should have, say, a 100x100x100x100 branching structure, or even bigger than that.

However, there is probably something much more clever going on here. The real aim is to determine the optimal strategy in the first step, where there are 128 branches, rather than in the final step of the tree. My question to the authors is this: does the chosen 128x16x2x2 structure give a reasonably accurate result? I am sure that there are, but in the time available to me for reading the paper I did not find the relevant references. However this looks like a very useful approach which is worth pursuing.

I was reminded here of an alternative approach to combining scenario paths with dynamic optimisation. Longstaff and Schwartz in 2001 looked at the pricing of American options — a problem which also requires dynamic optimisation. They proposed quite a different approach from the branching structure used in today's paper. What they did was to fit regression surfaces to the value function at each time step, working backwards from the terminal time. However, theirs was really a one-dimensional problem or a low dimensional problem, whereas today's paper deals with a multi-dimensional one. Perhaps the authors could comment on the relative merits of the two approaches.

For my final comments I come back to the main theme of the paper. The authors have described a global asset model with many assets. They then overlay a potentially complex dynamic optimisation problem and describe the various tools which exist to solve this type of problem. In the paper the authors note three types of error:

(1) sampling error - the result of using a finite number of scenarios

(2) parameter uncertainty - the result of having only a finite amount of historical data with which to estimate parameters

(3) model risk - any model we propose can only be an approximation to reality.

Suppose we have been able to eliminate the sampling error. Parameter and model risk are still quite significant and this led me to wonder if there is an element of spurious accuracy in the methods being proposed by today's authors. In order to illustrate my point, we can take a much simpler example, the one-period Markowitz portfolio theory model. The output from this model is well-known to be very sensitive to parameter uncertainty. In particular, if we make even quite small changes to the mean returns on different assets, then the optimal portfolio mix can change quite markedly. In today's paper we have many assets and many time steps so it seems likely that the consequences of parameter error might be at least as big.

On the other hand, this does not invalidate the optimisation exercise. For example, we may find the results of an optimisation exercise which are not, in fact optimal. Indeed, the resulting investment strategy might be quite different from the true optimum. However, in terms of expected utility to the investor we will often be very close to the true, maximum utility: that is, if we had known the correct model and parameter set. So even though the asset mix is very far away from the optimal asset mix, the actual effects in terms of optimal utility are small.

So the existence of parameter and model risk should not be used as an excuse for not trying to optimise expected utility or some other objective. However, work needs to be done to establish when, in the presence of parameter and model risk, dynamic optimalisation delivers a result where the actual utility value is indeed close to the optimum.

Reference:

LONGSTAFF, F.A., and SCHWARTZ, E.S. (2001) Valuing American options by simulation: a simple least-squares approach. Review of Financial Studies, 14(1):113-147.

Mr A. D. Smith: A few years ago, I thought dynamic optimisation was a prohibitively difficult problem. I looked at investment on the basis that a strategic asset allocation would stay in place for several years. We talked about stochastic optimisation, but it was a lot of effort to get it working and the answers made no sense. For example, in my 1996 paper, which the authors quote, I got enormous potential trading profits from the Wilkie model, without taking much risk. In the discussion of that paper, Professor Wilkie pointed out you could not rely on a model behaving like that and I agreed with him.

Judging from today's paper, much progress has been made. The authors are right to be proud of the generic dynamic optimisation tools they have developed. On this occasion they show an application to a utility problem. Dynamic optimisation problems also come up in a corporate context: for example, the management of profit sharing insurance funds. The same optimisation tools might be useful for the problems that occur here, too, although the objective functions would be rather different. If we believe this paper, then there is a whole industry building optimisation tools, which actuaries should be able to tap into.

As it stands, the process seems demanding on input and parsimonious on output. For example, I would like to see some in-sample testing. How much reward does the model say is achievable for a given constraint on risk? How does this compare to the out-of-sample tests of what is actually achieved?

Out there in the market, there might (or might not) be some gold, that is, some exploitable opportunities to make investment profits with low risk. If you want to mine these gold nuggets, you need to make sure any model correctly captures them. But the process of fitting a model, can generate what I call ‘fools’ gold’. By this I mean trading opportunities that exist only in the model and not in the real world.

Now we start mining, using dynamic optimisation. We think we find lots of gold nuggets — the question is: are they real gold or are they fools’ gold?

Calibration by statistical time series analysis, which the authors advocate in ¶3.1.1, or even by eye, which they also seem to advocate, is notorious for producing fools’ gold in simulation models. Parameter estimates might be unbiased individually, but the optimisation process serves to concentrate any sampling error. You can see this at work as follows. Simulate from your favourite efficient market model — so you know there is no real gold — then use one of your efficient simulations and apply the methodology in this paper to calibrate a second time series model. You will find that gold nuggets pop up all over the place in the second model, but you know they are not real because of the way you put the experiment together.

The paper is too dismissive, in ¶3.1.2, of economically based models. These more theoretically based models use asset pricing theory, which minimises the amount of fools’ gold. See for example Chapter 16 of Cochrane's 2001 book for a clear explanation of how this works.

The backtests in Section 7, which we have seen amplified on the screen, claim a performance that is impressive without being extravagant. As the authors point out, in ¶7.3.6, their moderate success is probably due to the position limits they impose, and in particular the constraint of turnover to 15% of the portfolio. This is like telling a gold prospector he can only take 15% of any nugget he finds, irrespective of whether its real gold or fools' gold. The total volume of gold that the authors declare is perhaps within believable limits, but in the small print they are asking us to believe that there is another six times as much out there and up for grabs. At this point I believe that greater scepticism is in order.

In conclusion, we should congratulate the authors for their ambitious steps forward in stochastic optimisation. These techniques raise the hurdles for asset models. For the current paper, the asset model is surely the weakest link. But as the old adage goes, the best way to get rich in a gold rush is to sell shovels. Then you do not care whether the gold is real or not. And you cannot fail to be impressed with the shovels in this paper.

Mr C. A. Speed. F.F.A.: The authors bring to the attention of the profession dynamic optimisation techniques, which have great potential for many areas of actuarial work.

In the pensions arena the work before us offers new possibilities for Defined Contribution (DC) schemes. The techniques in the paper could be applied to DC asset allocation strategies by setting investment strategies, which would take account of the utility of the member. We need to be cautious as there is much work still to be done in finding plausible utility functions and to be able to consider all aspects of the member's wealth and also personal liabilities. Then we have the problem of communicating that.

In the realm of Defined Benefit (DB) schemes, the paper has less to offer. The results make it clear that risk must be taken with investments when applying these methods. For any organisation there is clearly a limit on the amount of risk that can be assumed, particularly if that organisation might have need to raise further finance at a later date. So where should an organisation take risks?

There are two attractive possibilities.

(1) the company could invest in business projects in its current line of business. This should coincide with where the company has a competitive advantage.

(2) alternatively, the company could gear up its balance sheet and so use a tax advantage available for financing through debt instead of equity.

Both these approaches to risk-taking have clear advantages to the shareholders.

In contrast, taking risk in a pension scheme, for example, through equity investment, provides a shareholder with neither tax advantages nor the benefits of competitive advantage. As the amount of risk an organisation can take it has practical limits, equity investment or other risky strategies in DB schemes is clearly sub-optimal. The introduction of dynamic optimisation does not change this analysis.

It is a cause for concern that standard financial results we are aware of are not recovered within the current framework.

In ¶2.2.3 the authors talk about utility theory. Utility theory should be applied to individuals, as we are all aware. In this context the authors use it as applying to a fund. Using utility theory for a collection of individuals is fraught with difficulty. Essentially, we are using utility to measure the benefit to members. The utility function optimised seems to depend upon the fund, but in a defined benefit scheme there is also the implicit guarantee which is given by the sponsor and ultimately the shareholders. This does not seem to be allowed for. True we could update and optimise different functions but we are going to need to take this a step further. We know there is a guarantee, what value should a member put upon that or alternatively, what is the cost to shareholders of this guarantee?

We really ought to be able to value and to say something about that guarantee. The current paper does not, but this is possible, because we have arbitrage-free models, which, for example, include deflators. However these are precisely the models, which the authors reject in ¶¶3.1.2 and 3.1.3, so the key question relating to DB schemes cannot be answered in the current framework.

Again, this is another example of the problems of using econometric models rather than arbitrage-free models. Arbitrage-free models are harder for advisers. There are no magic solutions, which crop up and make everyone a winner, though they are a useful reality check.

In summary, the paper offers promise of dynamic stochastic optimisation. As a profession, I hope that we go away and learn from the important techniques presented. There is much work to be done still. We ought to be aware of unjustifiable simplifications or possible wrong terms, be it the use of the asset model or the representation of the economic interests of the different stakeholders. If we are not careful, these simplifications could lead to illusory gains rather than tangible results.

Mrs S. Bridgeland. F.I.A.: As Chairman of the Finance and Investment Board of the profession at the moment, I thank the authors for this paper.

With regard to the case study, it is useful to have a concrete example of a project in an area, which is of particular relevance to the profession. There are significant challenges in dealing with Defined Contribution arrangements as opposed to Defined Benefit arrangements, which this paper helps us consider.

There are four main challenges.

Firstly, the choice of the asset model when modelling for a single individual rather than a group of individuals merits special attention. In a group you expect there to be some sharing of risks over time, which permits some approximations. Previous speakers have mentioned the comments made in ¶3.1.2 about the different sorts of asset models that are available when modelling for an individual. I favour a scientific theoretically accurate model rather than, what Professor Dempster might have described as, a more artistic model that takes into account some of the real world risks. However, the profession should not be focusing too much debate on whether a model is particularly right or not, when the real issue for the profession is how wrong it could be.

Our responsibility is to ensure that we do not advise individual members of pension schemes in a way which might mislead them about the potential risks of a particular investment strategy.

Second, it is vital to have a better understanding of utility functions. I agree that using a dynamic approach to the asset allocation problem helps when modelling realistically. We know when advising Defined Benefit pension plan trustees that they have different utility functions when the market has just gone down by 20% compared to when it has just gone up. They have different demands for risks and return and a different way of thinking about scheme assets and potential rewards.

Much of the existing work on utility theory highlights the difficulties in actually desiring somebody's unique utility function, their real trade-off between risk and return. Whatever investment questions you ask someone, it is clear that when they come down to make investment decisions, reality kicks in and some other rules apply. For example, what they read in the paper last night about which way markets have gone may influence their attitude. We might be deluding ourselves, it might be another fools' gold to believe that we can actually tune into that aspect of human behaviour in a way that means that we can optimise the solution through a single model.

I do not see any harm in selecting a tractable model for utility functions if that means that we can find an answer but we need to be sure that we understand how wrong that answer might be.

Thirdly, communication. There are some useful concepts in the paper, for example, the expected shortfall across scenarios in ¶5.6.2, and the severe downside linear penalty in ¶6.4.1, which may be difficult to explain to an individual member of a pension scheme. But there is potential to use these sorts of models to develop new ways of communicating risk, and the compromises involved in suboptimal activity and suboptimal investment strategies to members. For example, looking at the traditional lifestyle approaches that we might use at the moment, how suboptimal are those? How much better could we do?

That leads on to the final challenge which is to help design better investment products targeted more on the risk and return requirements of individuals. In ¶6.3.2 the relaxed fit/rebalancing rule model looks like a promising move for further investigation.

I hope that this paper will trigger off further research to enable us to help those that we advise understand the risks involved in saving for their future.

Mr M. Lamb (a visitor; Managing Director, Investment Banking, Dresdner Kleinwort Wasserstein)**:** I am an investment banker. This paper provides a very important framework for understanding the trade-offs between risk and reward, not just for the actuarial profession, but also for others working with it as well. It laid down a number of principles for optimising wealth, not just for institutions but also for the individual market as well.

The debate raises two worries. Firstly about communication, for the non-actuary and maybe for the actuary. As Ms Bridgeland mentioned, there was a communication issue or challenge, not just communicating this type of analysis to the profession, but also to the institution and to the people, who will need to rely on this type of analysis. There is a far greater challenge probably for the institutions, the providers operating in the life and pensions industry, to communicate the benefits of this type of analysis to individuals, who are entrusting their pension assets to the industry.

When there is fragile confidence in the industry, there is an issue about how the results are interpreted as well. One of the benefits of this type of analysis is that it does identify risks. What is important is that those risks are interpreted as well and that judgements are made around that.

That brings me to my second point in terms of confidence. I think Mr Smith also made this point. It is very important that this is only a framework. It cannot replace judgement. Certainly, investment managers, actuaries and boards, will be using this as an important tool to make decisions, to make judgements, but it would not necessarily eliminate experience in terms of making those very important and what are often binary judgements.

In terms of the opportunities, this type of analysis enables us more accurately to quantify risk and analyse that form of risk, and there is an opportunity here for the investment banking community to tailor products to address those types of risks. But generally, if those risks are more precisely quantified in order to develop cheaper and more effective pricing for products to address risk, that should be beneficial for all people who place their assets in pension funds.

The President (Mr J. Goford, F.I.A.): I think that we have had some considerable meat in the contributions that we have had so far. I wonder whether Mr Dempster would like to reply to what he has heard so far.

Professor M. A. H. Dempster, Hon.F.I.A. (responding)**:** There appear to be two principal criticisms, which are not new. The debate is not new between generating asset free models and using econometric models. Both of these, as Dr Cairns pointed out, can be fitted to data. The question is what do you get with the fit?

We are taking the naive approach that history tells us something about the future. History tells us something about the relatively near future and not very much about the long future. If we are looking at long horizon problems, which we are, 20-30 years, then of course we need some kind of arbitrage-freeness.

However, this debate about the two classes of asset return models is more subtle because to generate the real world probabilities from a deflator model one has to tune it to something. The literature is a bit thin on this as much of it is in terms of, for example, state prices, Arrow-Debrem securities, and so on, which I find very hard to see or to buy in the marketplace, actually. Therefore this turning is quite difficult. I think that Mr Smith, Mr Speed, and Dr Cairns, would admit that it is equally difficult to calibrate these kinds of models as it is for generally specified econometric models, at least without using econometric techniques.

There is one criticism I should like to address immediately. Mr Smith made much about the fact that optimisers find the spurious things and then show you fools' gold, etc. This point is fairly easily addressed. One of our co-authors has been working on relatively simple schemes to make sure that those kinds of trivial arbitrages are not found and can be automatically eliminated as scenario trees are generated.

We were accused perhaps of being artistic, and that may be a compliment. We are also scientific in that we have taken a kind of schoolboy approach to spurious arbitrages which is to generate problems and see the outcomes. It was a surprise to me that suppressing the sampling error of the generated scenarios actually stabilises the decisions and, answers the question that Mr Smith raised: what happens in a sample?

Since the paper's topics were defined we have been spending our time studying the decision stability question. For example, Mrs Bridgeland referred to the probabilistic constraints and how these are related to individuals. It is critical that in assessing the scenarios that are generated to optimise the probabilities of violating a guarantee or a value at risk number, either of which might be considered, must use out-of-sample for flat scenarios generated exactly as the ones used in the model's scenario tree.

That turns out to be not as trivial as one might think, but with proper manipulation of parameters of the model one can ensure virtually the same probability of violation of a guarantee constraint within the model that has been optimised on, say, 10,000 scenarios as for 100,000 further flat scenarios that have been generated over the same time period from your scenario generator. That is a key result.

That leads me to a second misconception. Those of us who have been trained as economists, or in a latter day way, have come to financial economics, think of utility as applying to individuals. We have heard several comments along these lines. However, to go back to von Neumann and Morganstern a utility is simply a representation of preferences over probability distributions, and the probability distributions that count in the models of the paper are wealth distributions at particularly the horizon and possibly at time points before the horizon. As Mr Smith pointed out, as we roll forward we are only going to implement the first decision, the first portfolio balance, in the light of all the uncertainties. We want to handle the risk of the wealth distributions that come from the model. That is all we are attributing to utilities. You could call them objective functions, which we are using to shape distribution at the horizon or as we go along, which is more difficult, but it still can be done. There what we do is assign a utility function at each period. More recent experiments, than those in the paper, have been carried out by Michael Villaverde along those lines, although some of the early work is in the paper.

So I should like to dispel the notion that we are really exactly trying to get individuals. Mr Germano can clarify because he said in his opening remarks that we are dealing in fund design with a whole lot of different horizons. That is what is going to be needed to be provided to a bunch of individuals.

Another point, which was mentioned, is that we need to consider liabilities, as we go along. This is what the system was designed for, to deal with cash requirements or cash inflows as one goes over time and to see the possible nature of that over future scenarios, and their implications for today's balance for the fund.

Another point, made by Mr Lamb, was that no matter how complex the model, it does not replace human judgement. We used the term “decision support system” in the paper, and we meant exactly that. I tried to stress that with the graphics, which are a real technical challenge from a computer science point view, you want to be able to look at data that went into the statistical models, scenarios that are generated, and paths of the portfolios on the dynamic models to see what you would do in various scenarios that have been projected forward, and you would want to look at all that together and be able to flip back and forth between them. That is because the only way to get this right is to apply human judgement and knowledge. We are certainly not suggesting that this will be a turn on operation that just flicks out what you want.

Dr Cairns mentioned the yield curve model. We got the model from Lehman Brothers. It was one widely used for pricing fixed income derivatives in the past. It is a bit confusing, partly because we could not manipulate one symbol on the diagram; I hope we will solve this problem before the paper is published. (We have.)

It is essentially a three factor model. There is a very short rate, which is a one month rate, a three month rate, which is there implicitly because it defines is a slope with a long-term rate, which in this case is a 30 year rate, the Bellweather bond rate. Dividing up the yield curve this way, we are aiming to back out the market prices of risk for the three factors, not so much because we wanted to interpret them carefully but, because we knew they changed over time, even fitting them to today's yield curve data, and we wanted to see whether the variation of these quantities over time really correlated well with our macro economic variables.

Mr M. Germano: Starting from the last contribution, and the effort that has been put together in terms of communication, I will try to give more insight about wealth distribution and the study on the true distribution of the population in terms of true needs. We have also set up, in parallel with this working group, a research group for the past year. This will be a long-term partnership, in order to understand a bit more in-depth expectation of the client and how to interpret and to study the distribution of the population. So we are not trying to go down to the single individual future pensioner, but we are trying to classify them in a few classes and trying to answer with proper products. We are going close to a customisation of the products without going into the individual long-term plans.

This increases the number of inputs. The demand for new inputs is even higher, because we have a lot of work to do in order to understand what are the proper products to meet individual needs.

The true goal of the system is risk management, studying future scenarios in order is to limit the downside. The system is also used in order to develop stress scenarios, to verify where the current products already in the markets are really reliable and in which different scenarios would produce the wealth distribution expected as they arise.

Another point about judgement, the system has been built in order to be used with judgement and it is not replace the decision-making of the portfolio manager, it is just giving the portfolio manager an extra tool in order to analyse extreme scenarios and have more insight about the possible near and long-term future.

In terms of communication, a big effort has to be made. The client does not need to have all the information about the technicalities. What is important is to give a clear understanding of possible future scenarios.

Professor M. A. H. Dempster, Hon.F.I.A.: This paper is part of the effort of communication of this technology, which is being considered in the actuarial world, in the investment-banking world, and financial services generally. Just as in the early 1970s people began, 20 years after it was proposed, to look at the Markowitz model, people are beginning to look at stochastic optimisation models, 20 years after they were first studied. In the introduction to our paper we discussed the very restrictive assumptions under which myopic models are optimal. The inappropriateness of these assumptions and the advance of the technology generally is what is driving a lot of institutions, banks, actuarial consultants, and so on, to use some version of these kinds of models, and to try to optimise them.

Communication is a problem because they are complex. Regarding use for individuals, there have been several attempts round the world, by various financial services companies, to produce individual advice based on these kinds of models or approximations to them. These so far have been uniform failures. At least the banking industry is now backing off and seeing that so-called high net worth individuals are a more sensible market. There are several projects underway along those lines. These should come ultimately to individuals but not yet. Then utility, in the classical sense, will be usable. It is not, however, how we use it here.

Mr M. H. D. Kemp: I want to share one or two thoughts about asset liability modelling more generally.

When I have reviewed the results of asset liability modelling exercises, they gave the same relatively straightforward answer. The more you invest in equities, if you believe equities are going to deliver good returns in the long-term, then the better off generally you are, except in a relatively few circumstances, when you are worse off. This seems to be an almost universal result.

It seems, subject to a few sweeping generalisations, that with all such exercises you start off with a stance that is very low-risk indeed, and then you consider alternatives that move away from that stance in some shape or form. Whatever financial institution that an actuary might advise, the starting point is essentially to approach a very highly rated bank, or institution, and say: "Please sell me an instrument that will deliver exactly what my liabilities look like." In principle, let us assume that it is possible to do this. The next step is then to say: "This instrument costs 100, I have 120". Then there is a margin to play with. Conversely "I have only 80!" Then there is a need to find more funds. You always start with the low-risk position and then take a view as to whether and how to move away from that stance. "Arbitrage-free" or "not arbitrage-free" instruments and several of the tools covered in the paper can be used to produce possible outcomes.

Firstly, there is the problem of how to price a complicated derivative of the sort just highlighted that is to match exactly a specified liability. The methodology actually used in the marketplace seems to be very similar to the sorts of technique described in the paper. The techniques are widely used in derivatives houses for pricing complicated derivatives. Second, this approach also helps with the arbitrage-free vs not arbitrage-free question. A derivative house pricing such a derivative will work out some kind of low risk investment strategy or some other kind of dynamic hedging approach "to match" the liabilities. Therefore all the other aspects of the asset liability problem can be thought of as a divergence from the matched stance. If inefficiencies or arbitrages exist in an asset-liability model, then it is possible to work out where they are coming from by analysing the difference between the strategy the model claims is optimal when incorporating differential returns between asset classes and what it would have claimed as optimal if you had adopted a very high preference towards the lowest possible risk. In summary, there is a very strong link between some of the techniques described in this paper and some of the techniques already used by the investment-banking community, when pricing derivatives.

Mr P. J. Nowell: I have been using the techniques described in the paper for about 10 years or so now, that is asset-liability techniques, and mainly using them with with-profits funds. Therefore I should like to make a few observations.

The main thrust of the paper is the concept of optimisation. In my work I find it difficult, despite trying different ploys and interrogating the data, to hone in on the things that would help lead to optimal solutions and funding methods to do that in some way would be helpful.

My interpretation of history would be that in the 1990s, when we were looking at these issues, we started off with a simple conundrum of maintaining solvency continuously, as we are supposed to do in a life company, and to optimise what the asset allocation should be in order effectively to maximise returns to policyholders and, by implication, if you are taking 10% of the profits for shareholders, then to optimise that as well.

We did a huge amount of trial and error. What we tried to do, in order to come to some sort of conclusion, was converge on solutions as opposed to just trying out something and then trying something else and spending hours not converging. But that process involved more sitting and thinking what to do next, and trying to eliminate silly things, rather than any scientific optimisation process.

Later we focused on not only the actual return and the total bonus, but also how much of it we could give by way of reversionary bonus. That put a new dynamic into it. We came up with different types of solutions, and a better understanding of the underlying dynamics of the model. We then examined or modelled the different types of shareholders' transfers, the difference between a fixed and a floating benchmark, and, generally, what was the optimal strategy. This work caused us to think in terms of not just the quantum but the shape of the distribution. We were also able to use the model to calculate the cost of the support loans when acquiring other life companies. We learnt a tremendous amount about the dynamics of the funds but never found a clear way of optimising strategies.

I agree with the authors about the importance of use of graphics as a communication tool. The more ways you can turn the results around and zero in on particular issues, the better. We collect a tremendous amount of data, that is not always analysed fully, but to be able to look at it when necessary is very valuable. For example, the recent work of the Continuous Mortality Investigation Bureau, analysing the cohort effect, which was there all the time but which we had never seen before. Use of graphics is an extremely good way of seeing with great clarity something that is very difficult.

Adopting the models set out in this paper will enable us to get away from trial and error and come up with optimal solutions. My concerns are about the difference between arbitrage-free models and ones that are not arbitrage-free, the question to the authors is: do these techniques work well with arbitrage-free models? In other words, do they actually help you in the zeroing in on the right solution process, or is the optimisation heavily influenced by the fools' gold type of approach, in which case that may help to guide you in the right direction only if the model is correct but in which case you then have to worry about whether you are looking at what the model is telling you rather than what the real-world does.

Professor D. Blake (a visitor; Professor of Financial Economics, Birkbeck College, University of London and Director of the Pensions Institute)**:** When I entered my profession in the early nineteen seventies, the big thing was large-scale macroeconomic models. They started pretty small, beginning with a few equations, and then hundreds of equations and then thousands of equations. You could end up getting big government grants for building large-scale macroeconomic models of the U.K., the U.S.A., Japan, and even the whole world. This was fine until you started to predict. There were many little variations and permutations on this model. One of the variations always gave very good predictions. The problem was that you never ever knew which one.

After a while, research money dried up. It became clear that these large-scale models were of little practical use. They had too much of a 'black box' feel about them. This is also my feeling about the techniques used in this paper despite their sophistication. Instead of large-scale models we reverted to simple models, going back to first principles. I later became interested in pensions. My simple approach to this was the life cycle model, in which you want to smooth consumption over your life cycle. You have fluctuating wealth, fluctuating income, and so forth, but want to smooth consumption over your life cycle. A pension plan is just a way of smoothing consumption over your life cycle. It is an investment vehicle for switching resources from when you are young to when you are old, that is an investment vehicle to hold stores of wealth while you are young to accumulate a fund or funds from which you can drawdown when you are retired.

The key issue is the simple design of investment products, not the complicated models that we now see in the paper; models exhibiting all these risks, sampling risk, modelling risk, and so on. You can never be sure how precise and useful the predictions are, as was the case with large-scale macro models. I would always go back to first principles.

The models in the paper only focus on the accumulation phase. It does not discuss what to do during the drawdown decumulation phase. Ideally, two types of assets are needed, zero-coupon wage indexed bonds during the accumulation phase, to hedge the earnings risk during the working career, and during the decumulation phase something I call "survivor bonds" (see reference below).

Survivor bonds are annuity bonds where the coupons fall over time at the same rate as the cohort of the population drawing pension annuities dies out. For example, for every 100,000 people on the issue date of the bond, if, after a year, 98% of that group were still alive, then the coupon on the bonds would fall to 98% of the starting coupon. But the coupon payments would continue so long as cohort members were alive. Survivor bonds would enable life offices to hedge the mortality risks.

The interesting question, given all the arbitrage-free modelling in this paper, is why do the two types of assets not exist? Why do we need to look at emerging market funds and all the other asset categories discussed in the paper in order to find accumulation vehicles and decumulation vehicles for what is a very simple problem: transferring resources over time from when you are young to when you are old? The Pension Metrico approach to pension plans design is much simpler.

References:

David Blake and William Burrows (2001), "Survivor Bonds: Helping to Hedge Mortality Risk", Journal of Risk Insurance, 68, 339-348.

David Blake, Andrew J. G. Cairns, and Kevin Dowd (2001) "PensionMetrics: Stochastic Pension Plan Design and Value-at-Risk during the Accumulation Phase", Insurance: Mathematics and Economics, 29, 187-215.

David Blake, Andrew J. G. Cairns, and Kevin Dowd (2002), PensionMetrics 2: Stochastic Pension Plan design during the Decumulation Phase", Pensions Institute Discussion Paper PI-0103, Birkbeck College London (www.pensions-institute.org/wp/wp 0103.pdf)

Mr A. D. Smith: I should like to come back on the question of utility. Utility is a way of ranking distributions. That is where von Neumann and Morganstern start (reference). The question is what do you apply the utility to? von Neumann and Morganstern appear to be trying to apply it to somebody's total wealth. But what if you have a model describing only part of somebody's wealth? For example, perhaps somebody is buying a house and stock market investments. Maybe they have some bonds overseas and maybe some legacies in the pipeline from rich relatives. Does it make sense to take just the equities that they are investing in and put them through some sort of utility or Markowitz model? No, it would be better to get hold of the whole of their wealth to use that sort of model.

In the case of some defined contribution pension funds, you might find that the vast majority of somebody's retirement savings were in those funds. In such cases that might be appropriate to use a utility approach. For the paper, the client is a corporation that operates these funds and provides guarantees to those funds. To try to advise that corporation on the basis of the utility function of the fund seems a bit nonsensical. For a start, the corporation does not receive any of the upside, because that goes to the members of the fund. They are interested only when there is an event, when they have to make up the guarantees. When you are dealing with a corporation, rather than considering the corporation to have a utility, why not consider the investors in that corporation and their utilities? Very few investors have only a single equity in their portfolios, one should probably suppose that those investors can diversify and have already done so.

That rather changes the problem. It means that you might want to look at options from an option pricing point view rather than from a utility maximisation point view. There is no justification for applying utilities to corporations. If you take the utility approach to its limits you find that no corporations should ever specialise. As well as running investment funds, the corporation should also start making automobiles, organise package deals to far-flung places, exploring for crude oil, and so on, because you are going to improve the utility function merely by diversification. Of course, that only makes sense if the person who invests in this particular institution cannot diversify on their own.

That is the point that both Mr Speed and I were seeking to make. From an individual investment point view, we can see that at least in principle it makes sense to say let us try to figure out what this individual's utility function is. From the point view of managing a corporation, whether it is an insurance company with a with-profits fund, or whether it is a company with a defined benefit pension scheme, it really does not make a lot of sense to take just the fund in isolation and to imagine that somehow it was going to serve the corporation or its shareholders by applying a utility function to what is a tiny proportion of somebody's total wealth.

Reference

von Neuman, John & Morgenstern, Oskar (1947). *The Theory of Games and Economic Behaviour*. Princeton University Press.

Mr J. P. Ryan, F.I.A.: I should like to cover the issue which Mr Smith just raised, because I have recently been exposed to pension funds from a rather different angle, that is trying to find market solutions to some of the risks that one finds. The marketplace for a number of risks is very different from what the profession in general and the funds are investing for. The reason for that is in many cases there is not the proper marketplace, frustrating the use of a utility function for the individual corporation. If you look at issues like longevity risk vis-à-vis the fund, you cannot simply diversify. There is not a market out there that actually does that. Capacity is exhausted.

You start getting into a capacity pricing type approach rather than a utility type approach, which means that you go then into some of the non-life type assets because you cannot diversify. That starts putting the risk loadings up on these things quite substantially. The parameter for operation type risks again comes very significantly into play and significantly increases the price over and above the outcomes of a lot of these optimisation type issues.

The other issue, which in terms of pension funds where utility theory does not come into play, is often a lot of this depends on the credit standing of the individual institution that is standing behind it, which for the pensioner or the employee means that he has his risk very heavily concentrated in that institution. Again, that needs to be reflected and priced, which seems to be allied to the issues that Mr Smith raised.

There are many interesting points in the paper in terms of techniques, but when it comes through to pricing some of these things, we need to follow some of these matters through further.

The President (Mr J. Goford, F.I.A.): One of the advantages of sitting here, and in no way this being my field, is that I can ask some naïve questions.

Those of you who have read my Presidential Address may or may not be aware that I advocated much more involvement of actuaries with financial economics, and indeed recommended that all actuaries should read certain chapters of Brealey and Myers. That was the book of the month for July.

The book of the month for November, which I am two-thirds of the way through, and hence the naïveté of the question, is “When Genius Failed”, which is the story of Long Term Capital Management's demise.

So, my first question is: would what you have done have informed Long Term Capital Management, in particular on the issue of the assumption of constant volatility? The answer may simply be "No", in which case you can say so.

The second question is looking at what employees are now being confronted with in their defined contribution schemes. They are being given access at their workplace to information on their choice of investment advice or investment funds and being shown the different expanding funnels of doubt depending on what particular funds they have chosen from low-risk to high risk. Employees are encouraged to ask for information like "I do not want an expected return; I want a return that is 85% likely to happen." Would a fund manager with these sort of techniques enable a narrower funnel of doubt for a particular investment choice and therefore, for a given propensity for wanting an 85% probability of pension, enable a more risky underlying investment?

Question 3 relates to derivatives. Would what you are talking about enable an organisation to have a reduced demand for derivatives or indeed would it reduce the cost? Would decisions, that are likely to be supported by what you are showing here, reduce the need for esoteric derivatives or perhaps reduce their costs?

Professor A. D. Wilkie, C.B.E., F.F.A., F.I.A. (closing the discussion)**:** I start by making a comment on the system, that we have in the actuarial profession, for organising sessional meetings, or strictly, Ordinary General Meetings, of the Institute. Years ago, the Institute used to have six sessional meetings a year. That included the Presidential Address every second year. Now we have about nine a year and that is too many. We do not have a queue of good papers waiting to be presented, as some other organisations do. Instead, we look around for papers, accept almost any volunteers and then put pressure on authors to meet a deadline. As a result, papers are often late; they receive very little scrutinising and certainly not proper academic refereeing. The quality and reputation of the *British Actuarial Journal* consequently suffers. I suggest that we have fewer formal sessional meetings and take up all the spaces with discussion sessions, as sometimes we do have, introduced by unrefereed notes, not necessarily all printed in the *British Actuarial Journal*, as this meeting will be, but making a distinction between the refereed papers and the unrefereed notes.

I, and colleagues, are producing a paper for the next Faculty sessional meeting in January, and we are running late. I am conceited enough to think it will be an important paper, and it would be a better paper if certain aspects of it could receive more investigation, than would be possible in the remaining few days left to finish it off.

The paper before us this evening has clearly suffered from this problem. It has not been available for long enough for most of us to study carefully, and the lack of time to consider it has been reflected in the rather few technical contributors that we have had in the discussion, although there have been plenty of practitioners making extremely useful points in general. The paper has been hastily printed, so one needs both the A5 photocopy and the PDF version to get all of it. It also suffers from a necessary defect, that is common to all academic papers, which is it relies heavily on references to other papers. This is normal and quite proper in the academic field, because academics usually have access to good libraries and ought to have time to look up other papers on their relevant subjects, as research is part of their job. But actuaries, apart from those who are also academics, usually do not have access to the same quality of library or the time to make themselves familiar with all the previous literature. Of course, it will be difficult for anybody to do so in the space of a few days.

Here I can say these things, because the paper this evening is a good one. Its subject is perhaps the most important one that actuaries face: the guidance over time of a financial institution. It is not referring specifically to a defined contribution pension fund. The principles apply to any financial institution.

Most of us drive cars and also navigate a car to where we wish to go. Guiding a ship or an aeroplane, navigating it to its destination, are also skills that many have. Sending space ships round the solar system is also possible for those who have the skill. However, guiding a financial institution to avoid disaster and with the best possible outcome seems to be a lot harder. Part of the reason for this is that spacecraft obey the laws of Newtonian mechanics, which are well understood. There is some stochastic uncertainty, but it is less than with financial systems and the error correction mechanisms are better formulated. Ships and aeroplanes are subject to uncertain winds and weather but, barring accidents, they normally reach their destinations. We even manage it with cars, though traffic delays may mean that it turns out, in retrospect, that we have chosen a sub-optimal route.

Financial institutions are harder because there is so much stochastic uncertainty and the laws, if one can call them that, that guide them are not understood all, as we can see from the discussion this evening. Further, the overall structure of all but the simplest financial institution is very complex, and this has defeated previous attempts to tackle the problem comprehensively.

The authors this evening have shown us a serious and complex system for tackling this problem, or at least part of it. A real life office can influence its sales of new business; it can adjust its bonus rates in response to changes in the financial environment and in order to keep it on a good track. A real defined benefit pension fund may be subject to surprises from changes in legislation or accounting requirements and may have to deal with bulk transfers in or out. But there is often some flexibility and discretion in awarding pension increases or improving benefits and in contribution rates. If all else fails, one can close the scheme to new entrants or even to future accrual, as some schemes recently have done. It would be too difficult to add these real world complexities to what is already a very complex problem.

The methodology of the authors derives from two sources, which we have heard something about already. The first was dynamic programing, which Dr Cairns described briefly, to optimise over time. The other approach starts with linear programing. Like dynamic programing, it is a common part of any course in Operational Research. A great many OR techniques are rather like actuarial techniques applied in different fields, and much OR methodology can be useful for actuaries.

Linear programing maximises a linear function in an n dimensional space, probably subject to linear constraints. It is deterministic; it is not a trivial problem if the dimensions are high. Quadratic programing, what is needed to solve the Markowitz portfolio selection problem, maximises a quadratic function subject to linear constraints, and so on. Stochastic programing then moves up a step. Instead of a deterministic objective function of fixed constraints, we may have uncertainty in any or all aspects of them. Then dynamic stochastic programing, which is what the authors are doing, combines the two strands. One is optimising over multiple time steps using a Bellman principle with each step as stochastic and that produces an enormously large linear programing problem, that the authors have developed methods to solve.

Most of their methods are described in other papers referred to in their paper, but some of the methods may be proprietary and therefore not fully described in the published material and that is a pity.

As Dr Cairns said, to solve the problem one would really like a tree simulation with 1,000 steps on each branch or 100 steps on each branch so that there would be 100 choices in period 1, followed by 100 more in period 2, making 10,000, followed by 100 more in period 3, making 1 million. It is possible to specify such a problem, but it is just as impractical as to specify every possible game for chess, which can nevertheless be shown by logical argument to be a wholly deterministic game if only one had big enough computers.

In solving this problem, there are some things that the authors have mentioned at the very beginning. They talk about guarantees but do not mention, although mentioned in the discussion, option pricing methodology, and the comparable dynamic hedging that goes with that. The option pricing method determines a possible investment strategy with dynamic hedging investment, which requires pretty frequent portfolio revisions, and I wonder whether some splicing in of dynamic hedging in respect of the downside guarantees might be possible.

The econometric model in Section 3 gets a lot of emphasis in the paper. It deserves a separate paper, because the heart of the paper in Section 5 could be tackled with any particular model, not necessarily the one that the authors have presented. This model is very difficult to assess without knowing the parameters, and without seeing specimen outputs. There is a problem in putting in too many influences. It is not quite a vector autoregressive model, certainly not a linear one, but there are so many possible parameters, that you cannot necessarily see what the important effects are.

There has been discussion about the difference between econometric and arbitrage-free models. I think "arbitrage-free" is the wrong word. The Wilkie model, or any other of these types of models, is arbitrage-free in that there are no guaranteed profits to be made from particular strategies within the model, whereas they may be not market consistent in the sense that they may at times suggest that some particular investment has a very low rate of return as compared with others. I am a little surprised that Mr Smith is continuing to support, as he would call it, the arbitrage-free market efficient models as representing the real world. The Wilkie model, or those that might have followed it, would not have lost money in the last three years, whereas those who believe in the random walk model and therefore thought shares always did the same thing every year, at least had the same possibilities every year, might well have lost much. I can also possibly say that the Wilkie model had not made any money in the previous three years! It thought that share prices were too high in 1997 and it is still thinks they are a bit high in some countries.

Utility functions have come in for a bit of discussion. I am happy to follow Professor Dempster's approach, that utility functions are there to allow decision-makers to choose between different probability distributions. I do not think that it matters, as Mr Smith was bringing up in his last point, that the decision-makers do not have total information about the potential beneficiaries of the scheme. They have to make some sort of assumptions, the best assumptions that they can, about what the right level of security and the right level of pension increase is and what the right level of bonus would be for those beneficiaries, knowing that they are not dealing with the total wealth of the customers. On the other hand, they are the people who have to make decisions about alternative probability distributions and they need some strategy for taking them into account. The particular utility function that the authors have used, which adds up utilities over the years, I am not 100% happy with. I like the downside quadratic part. The sigma utility function over years does not take account of a possible trade-off between years. You may want a smooth consumption pattern over the years, but an individual may not mind too much whether he gets a bonus this year, when he can take a specially good holiday, or a bonus next year, when he can take that good holiday. You can trade off one against the other, and it is quite different to have the certainty of a bonus, but uncertainty about the timing of it as opposed to uncertainty about no bonuses, one bonus or two bonuses in successive years. So, taking into account correlation between years would be nice.

This is the sort of paper that I hope some younger actuaries will take account of. I hope they will have time to study it and to learn about the really very complicated mathematical process that is at the heart of Section 5. Dealing with very large systems is necessary for solving problems. Dealing with very large numbers of simulations is necessary. The problem about doing huge numbers of simulations is that you get such enormous results all hidden inside the computer, and having some visualisation system for looking at the results you need can be very useful.

In my forthcoming paper I was producing one graph, which told me not an awful lot if I plotted it in one way, but I plot the same data in a slightly different way and see something rather striking coming out of it. It is just a matter of sometimes being fortunate in choosing the right way to look at something visually. It can give you lessons, which simply the numerical statistics do not do. The emphasis on visualisation presentation is important.

I hope many people in future study this paper and the other papers referred to in it, learn from it and take up this exciting subject, which I first discovered something about about 10 years ago when I first met Professor Dempster.

Professor M. A. H. Dempster, Hon F.I.A. (in reply)**:** I will tackle the President's questions, some of the other questions that have been discussed, and some of the points that Professor Wilkie raised.

On the point of non-constant volatilities, we do have non-constant volatilities in our model. What we did was transform the model in such a way that it had constant volatilities in spite of the fact that we began with state-dependent volatilities. For the emerging markets returns in equities and bonds, these are strategic asset/liability problems, so we are dealing really with emerging market indices: we use Generalised Autoregressive Conditional Heteroscedastic (GARCH [1,1]) models and the now usual sort of way of dealing with non-constant volatilities. That was just a small comment in answer to the President's first question.

What the President next questioned concerning LTCM amounts to asking, does the use of these models help the manager to narrow the funnel of doubt? That is the whole point of the exercise. The proof of the pudding is in the eating. I believe that some pudding has been eaten, but there is more to come!

So far as "Will this reduce the cost or demand for derivatives?", if the funnel of doubt goes down, the need for derivatives goes down as well. So the answer to that question is one hopes so to a certain extent.

The point was made earlier that these techniques can be used to price derivatives, particularly if you study risk management problems for derivatives. We, and other people, have in fact done that, so our technology has been applied for that kind of technique.

A point that Professor Wilkie, Mr Smith and Dr Cairns made was that they would really like to see great big trees. We can solve problems with great big trees. In the experiments, I mentioned earlier, about stability we have been solving some trees, 100x100x100 and perhaps by another 100. But that means that we will not be able to study a dynamic problem, which might be over 20 years. We are not going to be able to rebalance every year or assume that we can solve a problem the size of the universe or something. We are not going to be able to do that. However, that is what the experiments are about, the stability of both the portfolio and actually the extreme scenarios in which the guarantees are violated. What this is really about is to try to see how to get the same results as for the huge problems that run for hours with appropriately chosen scenarios. We have a few more tricks up our sleeves for this, which we have not talked about in this paper. But we are working very hard on that problem because it is the key limitation.

Some of our points have been reinforced in this discussion. With others the jury is out, but we have been arguing back and forth about these questions for some time.

Mr M. Germano (in reply)**:** In reply to Professor Blake, we have been looking into ways to simulate wage inflation and trying to find the instruments that will allow us to add that risk, as well as mortality risk. It is under development.

Regarding Mr Smith's query, the point of view of the guarantor rather than the single individual, we are keeping that separate. The guarantor behind the product has his own model to add to the risk. We are looking to achieve a minimum wealth for the class of individual, not to add the risk for the corporation in terms of not meeting the target. There are two separate models. One is an underlying model system in a bank and that is different from the function that is in the model for the class of the individuals.

Professor M. A. H. Dempster, Hon.F.I.A.: To take up the comment of Professor Blake, it is really about a simpler world in which people work in corporations, firms or professions. They finished training, started work and then retired. As Mr Germano mentioned earlier on, in the modern world that is not necessarily the case. People change professions. They are laid off. They change countries, and so on. So there are a lot more liabilities in an individual's life cycle. One of the points of this kind of modelling, is that it can be applied to this. In the first instance, it will only be a practical and economic solution if it applies to people who have a great deal of money and are facing some strange liability, shall we say. But, ultimately, we are going to be able to walk into a financial institution and have something tailored for each individual, using this kind of technology. However that may be many years away.

The President (Mr J. Goford, F.I.A.): The great joy of this paper was Mr Germano's introduction, where he demonstrated some serious customer needs focus which is very close to my heart, recognising that employees do have different needs, saving for pensions, for legacy, and those not so risk-averse.

I express my sincere thanks and the thanks of all of us to the authors, the opener and the closer and those who participated in the discussion.

WRITTEN CONTRIBUTION

Mrs N. T. Ralston: It concerns me that we could be moving yet further into a world in which the senior managements, and boards, of insurance companies may well struggle to comprehend fully the basis on which their risks are taken. The same could apply to investment banks or others which use the sort of techniques outlined in the paper. This is a particularly insidious risk, as it is hard for boards to admit.

Professor M. A. H. Dempster, Hon F.I.A. (final comments in reply on behalf of the authors). A few short comments picking up questions raised but unanswered in the discussion.

Dr Cairns asked how this work relates to fast Monte Carlo the techniques for American and Bermudan derivative pricing recently introduced by Longstaff and Schwartz (*op.cit.*). The theoretical answer has been noted by both Dr Cairns himself and Professor Wilkie – their technique is a variant of Bellman’s dynamic programming backward recursion in time, while we are using a mathematical programming approach which considers decisions at all time points (and all scenarios) simultaneously. Practically, such optimal stopping or fixed time impulse control problems are conceptually simpler than the strategic ALM problems treated in the paper. In practice, however, numerical evaluation of, say, long dated cross currency swaps in three state variables is in our experience equally computationally intensive.

To partially address Dr Cairns’ question on the three types of possible errors in the methods of the paper, we have seen in our experiments to date that sampling error far overwhelms statistical parameter estimation error and that the robustness of current decisions due to the effects of the averaging of responses to future alternative scenarios tends to suppress the latter’s effects once sampling error has been controlled. Unfortunately, model error is possible with any approach.

Several commentators have asked for further details on our asset return models. We draw readers’ attention to the working papers, cited as forthcoming in the paper, which are now available (http://www-cfr.jims.cam.ac.uk).

Finally, Professor Wilkie mentioned the possibility of dynamic hedging in respect of the downside guarantees of defined contribution pension products. This is an area of no little importance which is being investigated in our current research.