A Simulation Based Approach to Asset Allocation Decisions

Philip Booth & Alen Ong

Department of Actuarial Science & Statistics

City University

Northampton Square

London EC1V 0HB

Telephone: 071-477 8478

Fax: 071-477 8572

Summary

Mean/variance methodology has been commonly used as a basis for making asset allocation decisions. Sherris (1992) demonstrated how this approach was really a special case of a more general utility maximisation problem. This paper intends to carry this idea further by applying numerical techniques to obtain the optimal asset allocation strategy, as well as incorporating explicit constraints into the selection problem.

Une simulation basée sur les décisions de répartition des avoirs

Philip Booth et Alen Ong
Département des sciences actuarielles et des statistiques
City University
Northampton Square
Londres EC1V 0HB
Téléphone : 071-477 8478
Télécopie : 071-477 8572

Résumé

La méthodologie de la moyenne et de la variance est communément utilisée pour base des décisions de répartition des avoirs. Sherris (1992) a démontré que cette approche constitue en fait un cas particulier de maximisation des utilités. Cet article se propose de poursuivre ce raisonnement en appliquant des techniques numériques pour obtenir une stratégie optimale de répartition des avoirs, et pour incorporer des contraintes explicites au problème de la sélection.

1. Introduction

1.1 Asset allocation decisions are increasingly being taken with the help of stochastic asset/liability modelling. A number of practical difficulties remain in the development of optimal asset allocation strategies for an insurance company. In this paper, by focusing on the assets side, the intention is to illustrate how some of the problems can be overcome. In the case of an insurance company, the decision process can be regarded as being a utility maximisation problem (where the interested parties are policyholders and shareholders, of a proprietary company), possibly subject to constraints such as probability of insolvency being below a certain level. Some of the papers written on these topics include Sherris (1992), Dinenis & Dickinson (1992), Krinsky (1985) and Stowe (1978).

1.2 In this paper, by looking at a personal investor, who has no specified liabilities, we concentrate on two particular aspects of the asset allocation decision making process, to produce results which are important in their own right. The investigation is best looked upon as the pursuit of an optimal asset allocation strategy by a personal investor (perhaps investing in a unit linked pension fund) who is facing a particular utility function but also has a number of constraints on the investment decision. Simulation and hill climbing iterative techniques are used to solve for the asset allocation strategy, for such an investor, who requires inflation protection, over short and long time periods.

2. The Nature of the Investment Decision

2.1 Consider a personal investor who wishes to invest a sum of money in unit-linked funds for a fixed period. This could, for example, reflect the process of an investor choosing the investment funds for a unit linked personal pension scheme. One can look at the asset allocation decision, either from the point of view of an investor choosing his own asset mix, or from the point of view of an insurance company, managing a mixed fund, determining the optimal strategy, for its customers. Traditionally, various factors could be said to have been taken into account in determining the asset mix which would suit the investor best. One factor will be the expected returns. Obviously the investor will wish to maximise the average return on his portfolio. This however will be subject to an acceptable level of risk.

2.2 Traditionally, risk has been defined in terms of variance of returns (Markowitz 1952). This has been the case explicitly, where the asset allocation decision has been seen as a trade-off of mean and variance of return. But it has also been the case implicitly, where certain types of utility functions have been used and assumptions made about the distribution of investment returns, which lead to the decision maker only being interested in the first two moments of the probability distribution of possible outcomes.

2.3 In this paper, a broader view is taken of the idea of risk. Firstly, it appears unnecessarily narrow to view risk purely in terms of variance of return, when other moments of a probability distribution may be more important. Secondly, it seems realistic to impose constraints on the investment policy. Such an approach adopts some of the concepts discussed by Clarkson (1989). In general, risk should not simply be expressed in terms of variance of returns but, rather, a measure of the effect due to any adverse consequences. Thus, a more general utility maximisation approach is adopted. By applying the appropriate utility function to the possible outcomes of the asset allocation decision, the *required* risk-return trade-off for any outcome should then be accounted for.

2.4 Thus investors will face a utility function, which takes a functional form and has parameters which reflect their subjective preferences. In this paper, utility maximisation of an investor, via asset allocation decisions is considered in the absence of any liabilities. The functional form of the utility function will be determined by the way in which risk is perceived by the particular investor.

2.5 The approach is thus general, not relying on assumptions about the distribution of investment returns or the restrictive assumption that the asset allocation decision is a mean/variance trade off. The process also imposes upon the investor constraints on the investment outcomes, taking it slightly outside the utility theory approach: it is felt that this best simulates the actual decisions taken, in practice. Such a constraint may be, for example, the desire to have a probability of making a negative real return of less than 0.05. Section 3 below, discusses how a stochastic investment model was used to generate the distribution of investment returns, to be incorporated in the asset allocation decision.

3. The Stochastic Investment model

3.1 There are two basic classes of investment model: econometric models and time series models. Their advantages and disadvantages are discussed in The Report on the Wilkie Stochastic Investment Model (JIA 1992). After careful consideration, a time series vector autoregressive (VAR) model was felt to be more suitable for our investigation, its main advantage being the greater number of asset categories the model encompassed. VAR models are formulated purely on a statistical basis meaning that they contain no significant econometric input. There are mixed feelings as to the relative merits of VAR models. However, the possibility of using a VAR model to estimate parameters using limited data was attractive.

3.2 In the above mentioned report, it was pointed out that, although VAR models were widely used in economics, they lack a regard to underlying economic theory and suffer limitations imposed by their linearity. It was also remarked that despite these criticisms, they are straightforward to fit and there is always some degree of uncertainty as to underlying economic theory in econometric models. The parameters in a VAR model are usually estimated using ordinary least squares and from an appropriate historic data set. Starting dates of the data set depend on the availability of indices and may vary across the different asset categories. Inflation, cash, gilts and UK equities could be based on figures going back several decades whereas for index-linked gilts, the data will only begin from the early 1980's.

3.3 In addition, there is also the problem of choosing the order of the autoregressive process. Time series analysts often make use of guides in solving this problem, normally by optimising a suitable function of the data. Some of these approaches as mentioned in Chatfield (1984) include Akaike's final prediction error, Akaike's information criterion and Parzen's autoregressive transfer function criterion. The limitations of these approaches are that they sometimes indicate too many parameters and assume the data to be normally distributed.

3.4 A VAR model can be used to simulate possible future investment returns with the distribution of future returns relying on past data. In order to check the adequacy of the model for our purposes, it was used to predict the return for each asset category one hundred times. Thus a distribution of investment returns could be derived using simulation methods. The resulting figures were converted into means and variances of rates of return for each asset type over each of the twenty years beginning at 1993. On considering the results, it was felt that some changes should be made to the model: the means and variances of return from property were rather unreasonable due to the poor data and this asset class was discarded for this investigation; irredeemable gilts were also dropped as it was felt that low/medium coupon long gilts adequately represented this investment type. The growth factor for overseas equities was increased to bring their average returns in line with that of UK equities as it was felt that, over the long term, the mean return from overseas equities should be about the same as that from U.K. equities and that the difference was due to data inadequacies. It seemed appropriate that these adjustments should be made, as they did not take us away from the VAR methodology in principle, whilst allowing the model to reflect reality more accurately.

3.5 After adjustment, the nominal expected returns (annualised) over twenty years from the asset categories ascended in the following order: cash, long gilts, medium gilts, short gilts, index-linked gilts, overseas equities and U.K. equities. The variance ascended in the order: long gilts, medium gilts, short gilts, cash, index-linked gilts, U.K. equities and overseas equities. Those results were intuitively reasonable. The VAR stochastic investment model was then used to determine optimal asset allocation strategies for different investors with different degrees of risk aversion.

4. Utility Maximisation

4.1 Consider an investor with an amount of personal wealth W, wishing to invest a sum A in a mix of unitised funds. The performance of the funds depends entirely on the performance of the underlying assets. Thus, the investor can be viewed to have invested in the asset classes directly. Using the following notation, define :

A = amount of initial assets

w_i = proportion invested in asset category i

j_i = rate of return earned over the period in asset i

R_i = $(1 + j_i)$

4.2 The asset categories could be numbered as follows :

1. cash
2. short term gilts
3. medium term gilts
4. long term gilts
5. index-linked gilts
6. UK equities
7. overseas equities

We can now define:

$$S = A(w_1 R_1 + ... + w_7 R_7)$$

The aim of the investor will be to invest in the various asset categories, in order to maximise the expected utility of S with respect to w_i's, subject to certain constraints.

$$\max_{w_i} \; E[\, U(S) \,]$$

This decision making process could be described as the investment optimisation process.

4.3 The conclusions which the investor will draw from the investment optimisation process will be highly sensitive to the utility function and it is important that an appropriate one is used. Let us first consider the desirable properties of a typical utility function. In general, a utility function should be monotonic increasing i.e. people tend to prefer more (money) to less. Utility theory also suggests that people are risk averse; this is accommodated if the second derivative is negative. The function should also be continuous.

4.4 In relation to the investment decision making process, it seems appropriate that there should be a measure of risk incorporated by means of a parameter which can be changed readily. Sherris (1992) achieved this by using the utility function :

$$U(S) = -\exp(-S/r)$$

where r is taken to be the measure of risk tolerance. The larger the value of r the more tolerant the investor is to risk. With this form of utility function, it can be seen that increasing the amount invested, A, thus increasing S, would have an identical effect on $U(S)$ as reducing the value of r by the same proportion. This is not unreasonable, as the more is invested in the fund for given values of W and r, the less tolerant the investor will be to risk, which is consistent with the idea of aversion to risk.

4.5 In essence, any investment decisions based on this utility function will depend on two components: the risk tolerance level, r and the proportion of initial wealth invested, A/W. So, it would be possible to have a group of investors with identical values of r, but making different investment decisions as each could be investing different proportions of their total wealth. In order to simplify the problem, the remainder of this paper will only be concerned with the effect of r for a *given* proportion of wealth invested. It should be said that, for a general utility function, the asset allocation decision will depend on the initial amount of wealth. If this is the case, then inter-temporal comparisons have to be used to compare the utility of the accumulated amount with the utility of the wealth given up. A further advantage of using the form of utility function described above is that the asset allocation decision, for a given degree of risk tolerance, only depends on the *proportion* of wealth invested. The utility maximisation investment strategy is the same for any two individuals giving up the same proportion of their wealth and depends only on the utility of the accumulated asset portfolio. In order to keep the decision making process consistent for varying levels of A and W, it seems logical to redefine S as:

$$S = \frac{A}{W} (w_1 R_1 + \ldots + w_7 R_7)$$

Thus the ultimate proceeds of the investment are taken to be a proportion of the initial amount.

4.6 From the point of view of the fund manager, the utility function should be independent of the size of the fund, as far as the decision making process is concerned. If the fund is wholly owned by the investors, the latter bearing all of the investment risk, then the concept of initial wealth does not apply to the fund managers. However, in the situation where guaranteed liabilities have to be paid by the insurance company, it is likely that at least part of the company's estate will be required to provide a margin of solvency above the value of the liabilities. It is in these circumstances where the proportion of initial wealth (or estate) would need to be considered.

4.7 In general, the measure of risk aversion is defined as:

$$\frac{-U''(S)}{U'(S)}$$

where $U'(S)$ and $U''(S)$ are the first and second derivative of $U(S)$ with respect to S. The utility function used in this investigation has a measure of risk aversion equal to $1/r$. It does not depend on S and hence is said to have constant risk aversion. An example of a utility function with diminishing risk aversion could be :

$$U(S) = (z-1)\exp(-S/r) - z\exp(-S/ra)$$

where a and z are constants with $a \neq 1$ and $0<z<1$.

4.8 Depending on the reason for the investment, the investor may have specific boundaries within which investment outcomes are considered acceptable. Thus, for example, an investor, investing over a long time period may regard returns below one percentage point above the rate of inflation as highly undesirable. It is just such an outlook, which is believed to reflect reality well, but which takes us away from the traditional mean/variance framework for optimising investment decisions. The undesirability of adverse outcomes could be reflected through altering the utility function. However, a more direct and economical way of doing so would be to specify a constraint on investment policy, such as specifying a minimum probability of obtaining a particular nominal or real return. For example, the constraint could be a specification of the largest acceptable probability of the return being less than a particular value, thus:

$$P(j < kg) < \alpha$$

where j is the nominal or real rate of return on the fund (depending on whether the constraint is specified in nominal or real terms), g is the rate of growth in the RPI and both k and α are constants. Another, more complex, type of constraint, could be a constraint on the probability of underperforming cash returns by a particular amount. A probability statement similar to that above could be used, but with g representing cash returns.

4.9 Mean/variance analysis is really a subset of the more general approach we are adopting. A mean/variance decision making framework would be most appropriate when investment returns are normally distributed or when the utility function is quadratic. In reality, equity returns tend to be skewed to the right. Third and higher moments do affect optimal asset allocations and should not be ignored. If options form part of the possible portfolio of assets, then clearly the distribution of returns could be skewed significantly. All aspects of the distribution can be accounted for if simulation is used to find the asset allocation strategies which maximise the investors utility. This flexibility extends to the use of any utility function and the inclusion of constraints.

4.10 In most situations, it would be useful that positive holdings of assets are held at all times and that no borrowing from outside the fund is allowed. Thus a further constraint is that we require:

$$w_i\text{'s} > 0 \qquad \text{and} \qquad \sum_i w_i = 1$$

5. Optimisation : An Analytical Approach

5.1 If a normal distribution of investment returns is assumed, the asset allocation problem can be solved analytically for an exponential utility function with no constraints. The solution becomes more complex, the more asset categories are used and here, the analytical solution for these asset types is derived. The problem may be written as follows :

$$\max_{w_i} \; E\,[\,-\exp(tS)\,]$$

where $t = -1/r$. This is equivalent to :

$$\min_{w_i} \; E\,[\,\exp(tS)\,]$$

5.2 The return from the fund will be a linear combination of the returns from asset types available, the total returns on the fund will therefore be normally distributed. If S is normally distributed with mean μ and variance σ^2, then the expression, $E\,[\,\exp(tS)\,]$, will be equal to the moment generating function of S. So the problem can then be reduced to:

$$\min_{w_i} \; \exp(\mu t + \sigma^2 t^2/2)$$

This function can be solved relatively easily if there are two asset types available but becomes more complex for more asset types. In the case with three asset types available:

$$\mu \;=\; \frac{A}{W}\,[\,w_1E_1 + w_2E_2 + w_3E_3\,]$$

$$\sigma^2 \;=\; (\frac{A}{W})^2\,[\,w_1{}^2V_1 + w_2{}^2V_2 + w_3{}^2V_3 + 2w_1w_2C_{12} + 2w_1w_3C_{13} + 2w_2w_3C_{23}\,]$$

E_i and V_i representing the mean and variance of returns from asset i and C_{ij} being the covariance of returns from assets i and j. By replacing w_3 with $(1\text{-}w_1\text{-}w_2)$, the expression $(\mu t + \sigma^2 t^2/2)$ can be differentiated with respect to w_1 and set equal to zero to give :

$$w_1 = \frac{(rW/A)(E_1 - E_3) + w_2[\, C_{13} + C_{23} - C_{12} - V_3\,] + (\, V_3 - C_{13}\,)}{V_1 + V_3 - 2C_{13}}$$

5.3 A similar expression for w_2 can be obtained in terms of w_1 and solved simultaneously. Even in the situation with three asset types, it can be seen that the problem is quite complex to solve, analytically. The solution obtained in this manner may also include negative values for w_i's which for most purposes is undesirable. Another drawback is the assumption regarding normality of the distribution of returns, which is needed in order to solve the problem and which may not be satisfactory, as has been pointed out above. With the more complex situations which could arise in the multiperiod case, an analytical approach is rarely feasible. This method would also not necessarily be appropriate for different types of utility function or if an analytical utility maximisation is to be subject to constraints.

5.4 A simulation approach can generally cope with all the difficulties mentioned above. It also has the advantage of being able to incorporate other constraints mentioned earlier, in addition to ensuring positive holdings of assets. Using this approach, individual scenarios can be considered which may be useful in making a better assessment of the consequences of adverse outcomes. This requires the use of stochastic investment modelling and simulation. The process is outlined below.

6. Simulation Based Approach to Investment Optimisation

6.1 The investment model can be used to generate a large number of future scenarios, say one thousand or more investment returns for each asset type. Each scenario could, for example, project the inflation rate, the return from cash and gross redemption yields on gilts, together with dividend yields and growth rates from equities, for each year during the next twenty years. Within each scenario, the figures are projected simultaneously for all asset types, rather than individually, to ensure that there is a meaningful correlation of returns from the different asset classes through time. The returns from real asset classes and inflation are interdependent. From our starting position, one thousand possible investment scenarios could be produced over the period concerned.

6.2 A scenario is generated by projecting the relevant variables over the next twenty years one thousand times. Thus the number of possible scenarios that can be studied, at any one time, will be one thousand. This is sometimes known as a path approach. An alternative would be to use a tree approach, where one thousand more scenarios branch out from each epoch or sub-period, producing 1000^{20} scenarios at the horizon date. Although the tree approach would be more realistic, it is not computationally feasible, particularly where correlation of investment returns in successive years is assumed. The path approach will not produce significantly different results from the tree approach, as long as sufficient scenarios are generated.

6.3 Having generated these future scenarios, any of the intermediate periods could be taken as the ultimate period. For example, if 10 years were chosen, it would then be possible to accumulate a sum of A for each scenario, over ten years, given the proportion invested in different asset classes. The chosen utility function would then be applied to each of the investment accumulations and the expected utility would be the mean utility of returns from the investment over the ten years. The optimal asset distribution is that which maximises expected utility. This can be found using numerical optimisation algorithms, including the so-called 'hill-climbing' algorithms used in this paper. Constraints may readily be included in the process.

6.4 'Hill-climbing' techniques in general, involve finding a local maximum from the given starting position. Hill-climbing processes are by no means infallible. In finding a local maximum we are in no way certain that the global maximum has been achieved. This is rarely a problem in optimisation such as that we have described, as the function we are trying to maximise tends in most circumstances to have only one maximum point within the prescribed boundary conditions. This can be checked by starting the 'hill-climbing' process from different initial positions although unconstrained optimisation rarely, if ever, yields more than one different solution. Difficulties do arise in the case where constraints are imposed, such as that of ensuring a minimum probability of obtaining negative real returns. Although every effort is made to ensure that the optimal solution found by this technique is indeed the asset allocation which provides the highest expected utility and satisfies these constraints, there is no guarantee that the global optimum has been reached. However, experience of using these methods suggests that the distribution determined using these techniques will tend to be very close to the optimal distribution. This may be confirmed by considering the global optimum without these constraints. The numerical technique is certainly adequate for most practical purposes.

7. Results of the Optimisation process

7.1 Using these techniques, various aspects of portfolio selection can be studied. One aspect could be the effect of the horizon date on the optimal asset allocation. In this example, the best asset mix for a person investing 10% of his wealth will be considered over one and twenty years. For each situation, optimal asset allocations will be determined for varying levels of risk tolerance. The values of the risk tolerance parameter r chosen for the one year investigation were 0.005, 0.01 and 0.02. The values of r for 0.001

and 0.1 are also included to demonstrate more extreme cases of risk averseness and risk tolerance. The optimisation was also carried out with constraints on allowable investment outcomes. The results are shown in tables 1(a) and 1(b), with the optimal asset mix for all five values of r given, correct to the nearest percentage. The real rate of return per annum, on the total fund, is given for the optimal asset mix expressed as a percentage. The probability of failing to meet the constraint of achieving a positive real return is also shown as well as the probability of the portfolio producing a return less than that of cash over the period.

7.2 Table 1 shows the optimal asset allocation strategies for an investor over one year. Investment returns are translated into purchasing power terms before applying the utility function. This situation is similar to that of a personal investor, investing in a personal pension plan, close to retirement. A number of interesting features of the optimal strategy can be noted :

(i) As the investor becomes more risk tolerant, larger proportions of assets are invested in equities. The investor who is risk averse favours index-linked gilts and cash. Given that the returns from assets and the income of the investor is measured in real terms, a result that the risk averse investor favours index-linked gilts seemed intuitively reasonable.

(ii) Overseas equities are not included in the optimal portfolio for any degree of risk tolerance. This is due to the variability of returns from this asset class caused by short term real currency risk and the volatility of many overseas investment markets.

(iii) The portfolios which are equity based have significant probabilities of negative returns.

(iv) No portfolios contain conventional gilt edged securities. This is due to their significant fluctuation of real returns, not compensated for by higher expected real returns.

(v) If the investor constrains the investment decision, rejecting those strategies which

risk tolerance	0.001	0.005	0.01	0.02	0.1
cash	23	0	0	0	0
short gilts	0	0	0	0	0
medium gilts	0	0	0	0	0
long gilts	0	0	0	0	0
index-linked gilts	77	96	85	65	0
UK equities	0	4	15	35	100
Overseas equities	0	0	0	0	0
Real return (% p.a.)	5.70	7.06	7.79	9.12	13.42
P(negative real return)	0.04	0.03	0.06	0.15	0.27
P(underperfoming cash)	0.12	0.10	0.08	0.16	0.29

table 1 (a). 1 year optimal allocations with no constraints

risk tolerance	0.001	0.005	0.01	0.02	0.1
cash	10	0	0	0	0
short gilts	0	0	0	0	0
medium gilts	0	0	0	0	0
long gilts	0	0	0	0	0
index-linked gilts	90	96	92	92	92
UK equities	0	4	8	8	8
Overseas equities	0	0	0	0	0
Real return (% p.a.)	6.31	7.06	7.34	7.34	7.34
P(negative real return)	0.03	0.03	0.03	0.03	0.03
P(underperfoming cash)	0.12	0.10	0.08	0.08	0.08

table 1 (b). 1 year optimal allocations with constraint : 3% max. probability of yielding negative real returns

risk tolerance	0.01	0.02	0.05	0.1	0.2
cash	0	0	0	0	0
short gilts	0	0	0	0	0
medium gilts	0	0	0	0	0
long gilts	0	0	0	0	0
index-linked gilts	89	82	63	29	0
UK equities	5	11	27	54	74
Overseas equities	6	7	10	17	26
Real return (% p.a.)	4.09	4.32	4.90	5.75	6.41
P(< 2.5% real return)	0.00	0.00	0.03	0.06	0.10
P(underperfoming cash)	0.04	0.03	0.01	0.00	0.01

table 2 (a). 20 year optimal allocations with no constraints

risk tolerance	0.01	0.02	0.05	0.1	0.2
cash	0	0	0	0	0
short gilts	0	0	0	0	0
medium gilts	0	0	0	0	0
long gilts	0	0	0	0	0
index-linked gilts	89	82	63	37	37
UK equities	5	11	27	48	48
Overseas equities	6	7	10	15	15
Real return (% p.a.)	4.09	4.32	4.90	5.56	5.56
P(< 2.5% real return)	0.00	0.00	0.03	0.05	0.05
P(underperfoming cash)	0.04	0.03	0.01	0.01	0.01

table 2 (b). 20 year optimal allocations with constraint : 5% max. probability of yielding negative real returns lead to a probability of achieving a negative real return of greater than 0.03, the equity based strategies are rejected; although some equities are included in the portfolio.

7.3 Table 2 shows the optimal asset allocation strategies for an investor over twenty years. Due to the fact that the returns were not discounted in any way, the values of r here are not comparable with those in the one year investigation. Many of the observations made about the results over one year also apply in the twenty year case but with some slight differences :

(i) Cash does appear in any of the in twenty year optimal asset allocations. This is reasonable as cash would not be expected to provide a good hedge against inflation in the long term.

(ii) Significant proportions of overseas equities now feature in the results reflecting their higher expected returns and greater stability in real terms over longer periods.

(iii) The relatively low probabilities of achieving a real rate of return of less than 2.5% demonstrate how most real asset categories are considerably more stable if held for long time periods.

7.4 Considering Table 2 in isolation, we can note the following points about the optimal asset allocation strategy :

(i) Once again, as the risk risk tolerance level increases, the proportion of funds invested in U.K. equities and overseas equities rises.

(ii) With a risk parameter of 0.1, a portfolio predominantly in U.K. equities but diversified between overseas equities and index-linked gilts is optimal.

(iii) If the investor is risk tolerant, the constrained investment strategy has a significantly greater proportion of assets in index-linked gilts.

(iv) The constrained investment strategy can have a greater probability of underperforming cash. This is because the constraint on the investment policy shifts the distribution to the left as well as making it more concentrated around the mean. The constraint used in table 2 (b) was to restrict the probability of achieving a real return of less than 2.5% to 0.05.

7.5 If the distribution of returns on the fund were normally distributed, then we would expect the optimal asset allocations as determined above to lie along the efficient frontier derived using mean-variance methodology. The constraints on the investment strategy and the fact that the distribution of investment returns is not normal, however, may shift the optimal asset allocation away from such a region.

8. Conclusion

8.1 In this study, we have found the optimal asset allocation for an investor with long and short time horizons for different degrees of risk tolerance, with and without constraints on investment policy. We have found that a risk tolerant investor, investing over a long time horizon will diversify between U.K and overseas equities and index-linked gilts. Overseas equities do not feature in optimal portfolios, if the investor is investing over a short time period. The investor will not necessarily be on the "efficient frontier" as defined in modern portfolio theory.

8.2 The simple case of the personal investor as described may be extended to a more complex situation, incorporating liabilities and optimising over more than one period. Simulation methods allow problems to be solved with fewer restrictive assumptions. There is also scope for incorporating guarantees on the minimum return from the fund, which is equivalent to purchasing put options on the assets held. The cost of obtaining guarantees by restricting investment policy can be compared with the cost from purchasing options.

9. Acknowledgements

The authors would like to thank Robert Chadburn (City University), Shane Deighton (Tillinghast) and Iain Allan (City University & Scottish Amicable) for providing help with practical aspects of investment modelling and Tillinghast for providing the software.

Thanks are also due to the Association of British Insurers and the Association of Consulting Actuaries for supporting this research.

References

CHATFIELD, C. (1984). *The Analysis of Time Series - An Introduction.* Chapman & Hall.

CLARKSON, R. S. (1989). *The Measurement of Investment Risk.* Journal of the Institute of Actuaries, 116, 127-178.

DINENIS, E. & DICKINSON, G. (1992). *Allowing for Liabilities of Life Insurance Companies in Mean-Variance Models.* Geneva Papers on Risk and Insurance, (September).

KRINSKY, I. (1985). *Mean-Variance Utility Functions, Flexible Functional Forms, and the Investment Behaviour of Canadian Life Insurers.* Journal of Risk and Insurance, Vol 52, No. 2, 241-268.

MARKOWITZ, H. M. (1952). *Portfolio Selection.* Journal of Finance, VII, No. 1.

Report on the Wilkie Stochastic Investment Model. Journal of the Institute of Actuaries, 119, II, 173-228.

SAMUELSON, P. A. (1969). *Lifetime Portfolio Selection by Dynamic Stochastic Programming.* The Review of Economics and Statistics, 51, 239-246.

SHERRIS, M. (1992). *Portfolio Selection and Matching : A Synthesis.* Journal of the Institute of Actuaries, 119, I, 87-105.

STOWE, J. D. (1978). *Life Insurance Company Portfolio Behaviour.* Journal of Risk and Insurance, Vol 45, No. 1, 431-447

WILKIE, A. D. (1986). *A Stochastic Investment Model for Actuarial Use.* Transactions of the Faculty of Actuaries, XXXIX, 341-403