
MODERN INVESTMENT MANAGEMENT

R C Urwin, United Kingdom

1. INTRODUCTION AND BACKGROUND

(1.1) This Paper discusses the way that approaches to investment management have altered in recent years and considers what further changes are likely to occur in the future. The problem is viewed from the point of view of the supervisors of funds (eg, the trustees of pension funds), rather than the investment managers themselves. Thus, we are concerned with how to take the key structural decisions which can then lead to the most effective investment management.

(1.2) The Paper has concentrated on pension fund investment. This is a particularly interesting area to consider because of its relatively high degree of sophistication, illustrated by recent developments in the application of asset liability modelling techniques and the increasing use of more structured and quantitative approaches. Practitioners in the pensions field are likely to be at the fore-front of the majority of developments in investment generally.

(1.3) Modern investment management for pension funds has developed into a three-stage process in which decisions are made from the top down:

(1) Policy Asset Allocation

The long-term average or norm for asset allocation between the major investment classes (to a UK fund this would be UK equities, overseas equities, property, index-linked securities, fixed interest, cash);

(2) Tactical Asset Allocation

The short-term divergence from this policy driven by current market opportunities;

(3) Selection of Securities

The selection of individual securities within the investment class.

(1.4) While many trustees or investment managers would not necessarily agree that their process is formally constructed along these lines, it has been increasingly apparent that this has become the de facto standard in investment management practices.

(1.5) This Paper considers each of these areas in turn, giving the greatest attention to the subject of policy asset allocation. To build a sound foundation to this area requires the development of an understanding of investment objectives which is provided in the following section.

2. INVESTMENT OBJECTIVES

(2.1) Investment objectives for UK pension funds have evolved from general forms like the following:

- (1) ensuring that sufficient assets are available to meet liabilities as they fall due (this might be termed the 'fiduciary' objective);
- (2) maximising the return on the fund at acceptable levels of risk (this might be termed the 'financing' objective);

to a more explicit form typically along the following lines:

- (3) achieving a return that ranks well with other pension funds over three year periods (as the objective is general to any fund, I term this a 'generic' objective).

The period of three years has become established as a satisfactory length of time to judge the investment manager's effectiveness.

(2.2) Generic investment objectives similar to (3) above have been adopted by the vast majority of UK pension funds. There are several clear advantages to adopting generic objectives:

- they are simple and straightforward;
- they provide managers with an asset allocation guideline (they need simply to know the average asset allocation);
- results can be monitored easily relative to these objectives.

The major drawback is that such an objective loses the link between assets and liabilities that was loosely identified in objectives (1) and (2). This gives the objective questionable relevance. For example, it would be quite possible for this objective to be achieved but for the fundamental objective in (1) to fail to be achieved.

(2.3) This is not to say that adopting such a generic objective is necessarily wrong. However, for such an objective to be suitable for a particular fund, the trustees are implicitly adopting one of these three arguments:

(i) The 'first principles' argument

The industry average asset allocation is a suitable policy for their fund;

(ii) The 'second principles' argument

Their pension fund is quite 'average' and they are satisfied that the industry average asset allocation is a suitable policy for such a fund;

(iii) The 'practical' argument

The industry average asset allocation may not be the best one for the fund but it is sufficiently close to to be an acceptable policy to adopt.

Further exploration of how any of these arguments might be arrived at is deferred to Section 3.

(2.4) Frequently, trustees adopt the generic benchmark as a means of trying to avoid taking a policy decision. However, in reality the decision is taken by default. In the absence of any other steer, the investment manager will take the industry average position as his benchmark. Any other course of action would represent an unacceptable risk to him and (he would argue) his client.

(2.5) The next key issue is whether all pension funds can be adequately served by generic objectives. Given the wide disparity in the financial characteristics of pension funds (as regards maturity and funding level in particular), the adoption of the industry average asset allocation as the common policy for all would involve pension funds taking widely different risk profiles. Put another way, if pension funds tried to adopt a common risk profile in the form of a universally accepted level of prudence, this would then predicate asset allocation policy varying from fund to fund as a result.

(2.6) Given these points, we must expect a significant minority of pension funds to take a stance different from any of the arguments of (i) to (iii) in (2.3) above. In such situations the key questions are: what should the objectives be and what should the asset allocation policy be?

(2.7) Because the pension fund that takes this distinctive position has financial characteristics different from the majority of pension funds, it is essential that the objective should reflect these differences. All investment objectives ultimately must comprise a balance between the competing priorities of the desire to increase return and to decrease risk. This means that the pension fund's 'special' objectives must be drafted along these lines:

(i) maximising the expected future return on the fund (equivalent to minimising the expected cost of the fund);

(ii) limiting the possible variations of future cost, particularly as regards the possible worst case positions.

(2.8) In some situations it may be appropriate to introduce to this definition a third wish similar to the following:

(iii) ranking well with other pension funds over three year periods but accepting that the particular level of protection against risk that the fund is adopting will lead to divergence of performance relative to other funds with a generic approach.

(2.9) It is clear that the special objective has lost the attraction of explicitness that was the merit of the generic objective. However, it has gained financial relevance. In Section 3 we show the techniques that can be used to develop explicit special objectives. The link to achieve this is deciding on an asset allocation benchmark.

(2.10) By introducing an asset allocation benchmark we can provide explicit objectives for the investment manager and so provide a full framework for monitoring as follows:

- the manager's objective is to perform well relative to the benchmark return and the manager's performance should be monitored relative to this benchmark;
- the benchmark should be monitored by considering the benchmark return relative to movements in the liabilities to check how adequately it fulfils the special objective.

(2.11) We are now in a position to tackle the development of the benchmark.

3. DEVELOPING AN ASSET ALLOCATION POLICY

(3.1) There are two possible approaches to developing a long-term asset allocation benchmark policy for a specific pension scheme:

- asset liability profiling;
- asset liability modelling.

Asset Liability Profiling

(3.2) The basic principle behind asset liability profiling is that by building a detailed profile of a pension fund, it will be possible to establish a suitable asset allocation policy. To develop such a policy, data on the scheme needs to be scrutinised including the latest actuarial valuation and scheme accounts. However, it is likely that no new calculations will be required.

(3.3) Pension schemes can differ in three basic areas:

- **Liability Profile**
 - The maturity profile of the scheme (as illustrated by the balance between the value of benefits currently payable for pensioners and future liabilities for members in service);
 - The nature of the liabilities, in particular the nature of the link to inflation in the benefits payable;
- **Actuarial Position**
 - The current funding level of the scheme as measured by the relationship between the assets and the accrued liabilities.
 - The cashflow position i.e. whether the scheme has an inflow (having to invest income) or an outflow (having to sell assets to meet liabilities);

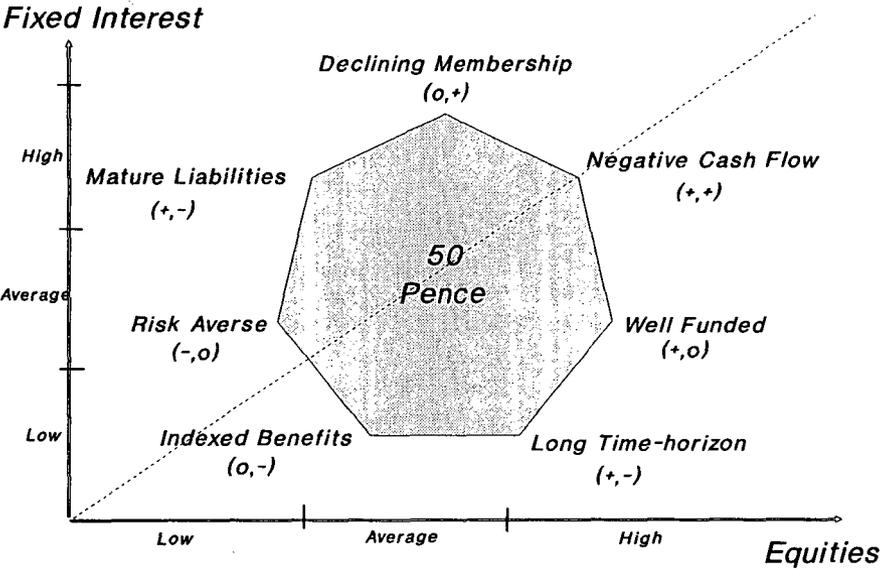
-
- **The Trustees' and Employer's Preferences**
 - The level of risk that the trustees and employer are prepared to accept;
 - The time horizon over which results are required to emerge.

(3.4) By examination of basic accounting and actuarial information it is possible to categorise a scheme according to each of these characteristics. The seven most common types of distinctive schemes are these:

- (1) the mature profile scheme - the majority of the liabilities relate to pensions currently in payment.
- (2) the indexed benefits scheme - pension increases are fully indexed.
- (3) the declining membership scheme - leavers and retirements exceed new entrants.
- (4) the well funded scheme - the scheme is in significant surplus.
- (5) the negative cash flow scheme - contributions plus investment income are less than the benefit outgo.
- (6) the risk averse scheme - the trustees and the employer take a very cautious view of the investment needs of the scheme.
- (7) the long time horizon scheme - an immature scheme in which the trustees and the employers have planned their investment policy over the very long term.

Pension schemes quite often fall in more than one category. Indeed, there have been instances where three or four of these characteristics applied.

(3.5) Research has suggested that each of these schemes has different requirements in relation to asset allocation. In relation to their weightings in equities and fixed interest - the two most important asset classes for a pension fund - we can illustrate below the most suitable mixes. For example, schemes with mature liabilities or declining membership generally require a higher fixed interest weighting than average; schemes that are well funded or with long time horizons generally require more than average in equities.



(3.6) Negative cashflow schemes should be orientated towards more liquid investments, higher income assets, and assets with growth in income. These characteristics suggest a preference for both fixed interest and UK equities and an aversion to property. Among the characteristics, we have found cashflow of less significance in most situations; the most important characteristics have been the maturity of the scheme and the funding position.

(3.7) This technique is termed asset liability profiling. The technique is obviously very crude but certainly provides a more structured discussion on investment policy and also is valuable in raising awareness of the characteristics of pension funds.

(3.8) The main difficulty with using the technique is that it requires a great deal of judgement to turn the profile into a specific asset allocation policy. It also involves comparisons with the average pension fund policy without being in a position to address whether this is a satisfactory starting point. Asset liability modelling tackles these difficulties.

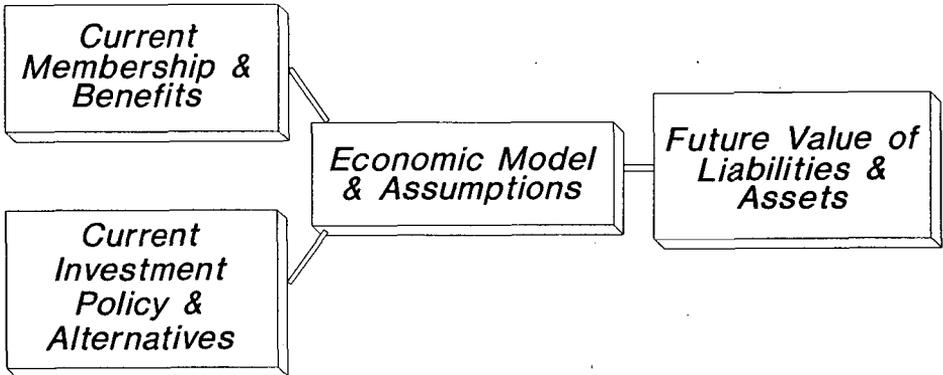
Asset Liability Modelling

(3.9) While the processes of asset liability modelling may be complicated, the principles of using modelling are straightforward. Modelling simply involves making projections as to the future financial position of a particular pension scheme. Calculations reveal the likely funding position of the scheme (the value of assets in relation to accrued liabilities) and the contribution rate at some future date, perhaps five, ten years or further into the future. The process is used to compare the respective merits of a range of asset allocation policies both as regards their rewards and their risk.

(3.10) The process is illustrated below. It is basically an extension to an actuarial valuation in which:

- the starting point is the current data as to the membership of the scheme and the current investment policy;
- the driving force is the economic model which calculates the future payments of benefits and the investment income and capital gains from the investment policy according to certain assumptions;

-
- the end result identifies the value of liabilities and assets at some future date.



(3.11) Generally, modelling systems calculate and compare the most likely position with the extremes in terms of worst case expectations (the 'downside') and the best case expectations (the 'upside'). This information is used to gauge the uncertainties implicit in any investment policy and therefore the risks that the employers and trustees actually face.

(3.12) Asset liability modelling is now practised through a wide variety of processes. The common characteristic of these processes is that they all involve the stochastic analysis of the future financial position of pension funds. It is usually but not always the case that modelling processes will also have the following characteristics:

- lognormal type distributions are assumed for investment return;
- time series analysis is used to identify asset liability relationships;

-
- assumptions for return and risk combine historic data and prospective judgement;
 - Monte Carlo techniques are used to generate results;
 - optimisation using quadratic programming techniques is used to identify asset allocations that are efficient (ie have minimum risk for any level of expected return).

(3.13) It is apparent from this specification that these processes are complicated. Models are developing in sophistication all the time and much further work remains to refine them.

(3.14) Despite the considerable research carried out on modelling, it remains a problematic process. The crucial difficulty remains that it is not possible to set assumptions with precision. Given the sensitivity of the results to variations in assumptions, the output from modelling must be regarded as 'fuzzy' (ie subject to error).

(3.15) The experience of modelling work to date suggests that this warning has not always been heeded. It is common to see modelling results presented as precise catalogues of future outcomes. It is also common for optimisation to be used slavishly and rigidly so that a choice of policy is taken exclusively from the set of efficient portfolios (the so called 'efficient frontier'). In reality, all modelling outputs are crude estimates and 'fuzzy' results should never be used in such an automatic fashion.

(3.16) One further problem with optimisation relates to the risk measure that one seeks to minimise. In most situations, there will be several measures of risk of relevance to the trustees. Optimisation can only deal with one. It is important to differentiate between the three types of application of the model and to emphasise the need to carry out each:

Type 1 - Testing; various asset allocation policies of relevance and interest to the fund (like the industry average asset allocation for example) are tested to check what future financial profile (eg future funding level) is likely or unlikely to result;

Type 2 - Optimisation; working from various future financial profiles (eg in terms of expected funding level), asset allocation policies are identified which minimise the possible risks (however defined);

Type 3 - Reverse Optimisation; working from various asset allocation policies, assumptions are developed which make such policies optimal (where optimal is defined as in Type 2 as a policy which for that level of probable reward, possible risk is minimised).

(3.17) It is clear that Type 2 applications are the most powerful and also the most dangerous. It is only with a full exploration of the issues through Type 1 and Type 3 applications, that Type 2 results can become useful.

(3.18) These points serve to underline the importance of modelling being the servant rather than the master of those using the process - the actuaries, consultants and trustees. Modelling provides a disciplined quantitative framework for qualitative discussion on asset allocation policy. It adds significantly to the quality of such a discussion. However, in no sense does it take away the need for this discussion. The best decisions on asset allocation policy will always involve subjective judgements.

(3.19) As a result of the discussions built around a modelling study, trustees can take the crucial step of deciding their asset allocation policy. Their decision might be that the industry average asset allocation is reasonable enough. Given the comments earlier about fuzziness, it is fair to suggest that this will apply to a significant number of funds which will satisfy themselves that the industry average asset allocation lies on or close to their (fuzzy) efficient frontier.

(3.20) Others will decide that their objectives are not adequately served by the industry average asset allocation and can decide through the above process what their asset allocation policy should be.

(3.21) In either case, trustees must now consider implementation.

4. INTEGRATING POLICY WITH TACTICS

(4.1) We return to the critical distinction between policy and tactics made first in Section 1. The development of policy is a joint effort in which the trustees and employer must ultimately take responsibility, relying heavily on advisers. When it comes to tactical decisions only the manager is in a position to take responsibility.

(4.2) Managers have become very familiar with the operation of funds with generic objectives. They readily take tactical asset allocation decisions by overweighting or underweighting an asset class on account of a short-term investment view. They also readily take stock selection positions.

(4.3) Some complication arises where managers themselves have a policy position relative to the industry average, commonly an overweighting in equities. The interpretation of this can be awkward. Whose responsibility is this - the trustees' responsibility in selecting the manager or the manager's responsibility in being accountable relative to the industry average? This point can only be considered on a case by case basis.

(4.4) Where an asset allocation benchmark has been specified, the managers must adapt their tactical asset allocation to reflect the benchmark. The tactical decisions taken must have the objective of adding value without compromising the underlying policy. This suggests that there are two conditions that should apply to the asset allocation actually adopted:

- the average asset allocation in time should be equal to the benchmark;
- the deviations from the benchmark should not be too great (an operating band can be introduced which encourages adherence to these parameters).

(4.5) While this appears straightforward, managers have limited experience of working to this system and often adapt their generic style inconsistently. This may be because the distinction between policy and tactics is often indistinct in the generic approach and therefore the tactical decisions which should technically be expressed simply as overweight and underweight proportions ('load differences') have not previously been explicitly identified.

(4.6) The following example might make this clearer. A manager with 70% in UK equities relative to the industry average of 60% might be expected to adopt the same overweighting of 10% in a special brief. Thus a benchmark of 50% in UK equities might lead to 60% as the current tactical position. However, the manager might take the view that the 10% overweighting in the generic brief was attributed in part to tactics, in part to policy. This problem can only be overcome when all parties explicitly identify their policy position, a key step forward in modern investment management.

(4.7) Further complications can arise when tactical positions for a generic brief cannot be implemented on a special brief as it would lead to less than 0% or more than 100% in certain asset categories. Some ad hoc solutions will be required, none of which will be complex.

(4.8) Risk was fully explored in the policy decision but must be considered again in tactical decisions. It is clear again that again a balance needs to be struck:

- on the one hand, assuming the manager has above average skill, you would wish to give more scope for his tactical decisions to add returns to the fund;
- on the other hand, you would not wish to allow too much latitude which would add to risk and would diminish the benefits from the policy decision.

(4.9) Risk in tactical decisions can be adequately measured by the volatility of results relative to the benchmark return. In practice for UK pension funds, this volatility (standard deviation) is generally in the range of 0% pa (passive management) up to 5% pa (very active management). In practice, the vast majority of managers in the UK would fall in the 2% pa to 3% pa range. At this level variations in performance (and therefore risk profile) would rely more on policy decisions than tactical decisions. In conclusion, such tactical decisions appear consistent and integrated with the policy provided they are likely to add value.

(4.10) Tactical decisions are in aggregate across the industry a negative sum game. Therefore in establishing the trustees' approach to tactical decisions, some consideration of passive approaches is worthwhile.

Passive Management

(4.11) The concept of passive management at the stock selection level is now widely understood. A passive fund incurs certain transaction costs which lead to a return that is likely to be slightly less than the return on the relevant index. Aside from this point, indexation methods in stock selection have proved they can generate results which exhibit very close performance tracking relative to the relevant index.

(4.12) The equivalent concept of passive management in tactical asset allocation has not yet been developed. However, the concept is equally valid in this context. In developing the long-term asset allocation policy, it would generally be assumed that a policy of say 50% in UK equities would stay at 50% (rather than varying with market movements). This former approach is often referred to as a constant mix policy in contrast to a buy and hold policy for the latter approach.

(4.13) With a constant mix benchmark, rebalancing needs to be carried out to return the allocation to the benchmark. Various decision rules might be used which would achieve this rebalancing without incurring too heavy a level of transaction costs or reducing performance.

5. MONITORING AND REVIEW

(5.1) Monitoring has to cover two distinct areas:

- **Policy:** progress of the policy relative to the liabilities;
- **Tactics:** progress of the managers relative to the (passive) benchmark.

(5.2) To achieve this separation, a calculation of the asset allocation benchmark return is required. This will depend on what form of policy is agreed. In the most common situation where a constant mix policy is used, the calculation must involve an implicit rebalancing of the fund back to the constant mix after market movements take the asset allocation away from the constant mix. The frequency of this rebalancing needs to be specified: this can be continuous or periodic, monthly, quarterly or annual. The choice must be consistent with the policy selected and the practical limitations that the manager is working within.

(5.3) The calculation of a continuously rebalanced benchmark involves an unfamiliar equation but is very straightforward.

For asset classes i , j , returns $R(i)$ and $R(j)$ benchmark proportions $P(i)$ and $P(j)$, the Benchmark return is

$$((1 + R(i)/100)^{P(i)} (1 + R(j)/100)^{P(j)} - 1) \times 100$$

(5.4) Benchmark returns can be built up from sector indices or sector medians (taken from performance surveys). Where the indices are well constructed and representative of the portfolios that fund managers construct, then the use of the index seems most appropriate. In other situations, the use of the sector median may be a more practical alternative.

(5.5) Comparing the manager's return with the benchmark return will show whether the manager's tactical decisions have increased performance. These results should be seen in the context of medium term objectives (three years or so). Comparing the manager's return with the industry average return in these situations is likely to be confusing.

(5.6) Monitoring of the policy should be relative to the liabilities. The performance of the policy must be considered both as regards returns and risk. So the closeness of tracking of the liabilities must be considered.

(5.7) This comparison must be seen in the context of the objectives which are set in the long-term (upwards of 10 years). This makes the process difficult as in most situations the policy will change more frequently than once every ten years. This makes the monitoring largely qualitative with quantification representing a check on the process.

6. CONCLUSION

(6.1) The evolution of investment management into a three stage process:

- Policy Asset Allocation
- Tactical Asset Allocation
- Selection of Securities

has been positive in bringing greater structure and discipline to investment decisions.

(6.2) Using this methodology it is possible to design investment objectives to be better integrated with the trustees' and employers' financial objectives.

(6.3) Asset liability modelling will play an increasing part in the development of such integrated policies.