

The Practicalities of Budgeting, Managing and Monitoring Investment Risk for Pension Funds

Portfolio Risk and Performance Working Party

Mike Brooks (chairman)

David Bowie

Martin Cumberworth

Alistair Haig

Bernie Nelson

**Faculty and Institute of Actuaries
Finance and Investment Conference**

Guernsey

24th-26th June 2001

Executive Summary

In this paper we discuss:

- the issues involved in establishing a risk budget for a pension fund
- how the active risk may be budgeted across asset classes
- what trustees should be looking for from their investment manager to ensure that the fund is being managed in line with the risk budget

Risk budgeting

The investment risk of a pension fund is a combination of strategic risk and active risk. Strategic risk is the risk of the strategic asset allocation of the fund relative to the fund's liabilities. Active risk is the risk taken by the investment manager relative to the strategic benchmark.

We have demonstrated that for the majority of pension funds strategic risk swamps active risk. This provides support for two of the recommendations from the government-sponsored Myners report:

- “The attention devoted to asset allocation decisions should fully reflect the contribution they can make to achieving the fund's investment objective.”
- “Where they believe active management to have the potential to achieve higher returns, funds should set both targets and risk controls which reflect this, allowing sufficient freedom for genuinely active management to occur”.

A significant complication is that, given the wide range of solvency measures that exist (e.g. on-going, buy-out, MFR/FRS17), the level of strategic risk is multi-dimensional. This makes it difficult to set (and confusing to understand) the level of strategic risk of a fund.

The financial economics perspective

Finance and economic theory would suggest that, to first order, strategy is irrelevant. The pension scheme is not an economic entity and so does not add economic value. It represents just part of shareholder wealth. A change in the risk and return profile of the pension scheme impacts on the contributions required to the scheme from the sponsoring company which in turn impacts on the risk and return profile of the company's profits. Although returns to shareholders may increase, these are accompanied by an increase in systematic risk.

Shareholders can and do rearrange their assets outside the pension fund in order to establish their own chosen levels of risk and hence trustees serve the shareholders poorly by agonising over how much or little systematic risk to take on. All that putting systematic risk in the fund does is complicate shareholders' lives when it comes to establishing their own personal levels of investment risk.

Although this viewpoint would seem to be rather gloomy for consultants and equity managers, it does have a silver lining. Because systematic risk is not fundamentally crucial to the pension fund, active risk comes back into its own.

If active managers can be found who are genuinely adding value by exploiting inefficiencies in the market or removing market frictions and if pension funds have an advantage over individuals in employing the managers then such investment by the pension fund might be adding economic value. Consequently there is also value in consultants helping pension funds to find such managers.

Even if the stakeholders of the pension scheme accept the above theory they may still prefer to take on systematic risk within the pension scheme in exchange for lower contribution rates. Although the financial economics perspective is now becoming mainstream within the actuarial profession, it remains to be seen whether the wider acceptance of the theory will result in wholesale changes in pension fund investment practice.

Allocating the risk budget within asset classes

Having set an overall active risk budget we have outlined a framework to consider the appropriate levels of risk to run within each asset class. The relationship between risk and return has been explored and a reasonable set of assumptions outlined from the perspective of a fund manager. In particular we have considered the assumptions to be made regarding levels of skill assumed across asset classes and how skill declines as the risk level is increased. Using such assumptions we have demonstrated how to set a risk budget for asset classes using an optimiser, seeking to maximise returns (or information ratio) for a given level of total active risk.

Risk Management and Monitoring

We discuss the issues that pension funds should consider in establishing risk objectives for their managers and what they should look for in managers to ensure that they have appropriate procedures in place in order to manage to these objectives. We advise against an obsessive focus on short-term active risk that may hinder the overall objective of enhancing long-term returns. Given the minimal impact that active risk has on the overall risk of the fund and the limitations in predicting future risk levels the risk objectives should be fairly broadly defined.

A sound portfolio risk management process will look at a range of measures. Whilst recognising the limitations of the tracking error measure we suggest that the fund uses broad tracking error ranges as the primary guide to the level of risk to be run by their manager(s). This should be supplemented by other measures (e.g. guidelines on maximum country, sector and stock positions) as required.

1. Introduction

In setting the investment policy of the fund the trustees of most pension schemes have the following objectives:

- To ensure the security of members (accrued) benefits on an on-going and discontinuance basis
- To produce a smoothed contribution rate for the sponsoring company
- To achieve high long-term investment returns and hence reduce the contributions required from the sponsoring company

The first two objectives encourage the trustees to follow a policy whereby the assets are closely matched with the liabilities. The final objective encourages a move away from this matched position into assets that are expected to generate higher returns over the long-term.

The scheme's investment policy therefore needs to reflect the extent to which the trustees are prepared to take on the risk of a less stable contribution rate and potentially less security for members' benefits in order to aim for higher returns and lower long-term contributions.

In practice a pension fund sets its investment policy in a two stage process:

- The trustees establish an appropriate asset mix for the fund. This is the strategic benchmark for the fund and outlines the balance between bonds, equities and property, and between domestic and international assets. The **strategic (or policy) risk** of the fund is the risk of poor performance of the strategic benchmark relative to the value of the liabilities.
- The tactical implementation of this strategy involves the selection of one or more managers (if not internally managed) and a decision on the appropriate level of risk that these managers should take *relative* to this strategic benchmark. This is commonly known as **active (or manager or implementation) risk**.

In recent years there has been an increased focus on this process. The term “risk budgeting” has been coined to refer to the process of establishing:

- (a) how much investment risk should be taken
- (b) where it is most efficient to take it in order to maximise returns

There has also been an increased focus on how investment managers go about managing the active risk of the fund.

In this paper we discuss:

- the issues involved in establishing a risk budget for a pension fund
- how the active risk may be budgeted across asset classes
- what trustees should be looking for from their investment manager to ensure that the fund is being managed in line with the risk budget

We are also keen to promote an understanding of the investment risk of a pension fund in its entirety. In particular we feel that the interaction between strategic risk and active risk is poorly understood. This misunderstanding, in conjunction with an overly short-term focus, can (and has) lead to a spurious and inappropriate emphasis on controlling the range of manager excess returns to lower levels than is desirable.

1.1 Interaction between strategic and active risk

It is widely accepted that it is possible to achieve long-term enhanced returns by moving away from the (theoretical) liability-matched investment position. However there is no free lunch in this strategic asset allocation decision. Any increase in projected long-term returns is accompanied by an increase in systematic risk (see section 2.2. for further discussion).

The key focus when setting the strategic asset allocation is therefore on the risk tolerance of the stakeholders of the scheme (i.e. trustees, scheme sponsor and members)¹. In other words how much systematic risk are they prepared to take on in an attempt to enhance long-term returns. In practice the behavioural biases of the decision-makers will also have a major bearing on the strategic asset allocation.

The decision on active risk is somewhat different. Active risk is typically independent of strategic risk. The level of active risk is also typically far smaller than the level of strategic risk. It is therefore largely diversified when combined with the strategic risk.

The key question on active risk is whether the trustees believe that active management generates positive excess returns, or more specifically whether they believe that, with the help of their advisors, they can select a manager (or managers) that generates a positive excess return. If they believe they can then the risk return trade-off for taking on active risk is very attractive. This is best illustrated by way of an example.

Example

In order to understand how the level of active risk impacts on solvency risk we need to consider the relationship between the active risk and the strategic risk. Urwin et al (2001) assume that the strategic risk tracking error is 13.6% and the outperformance generated from strategy is 2.0% p.a.. In section 3 of this paper we simulate how the excess return from the manager changes as the level of active risk increases.

Using these assumptions as our base, and assuming that active and strategic returns are independent of each other, we can then see how increasing the level of active risk impacts the overall risk-return profile of the fund:

¹ In theory this is purely the trustees decision but in practice given that it is the members that they are acting on behalf of and that it is the scheme sponsor who is (wholly or largely) bankrolling the scheme all parties will tend to have an influence in the decision.

Table 1: Understanding the interaction between strategic and active risk

Manager	TE	0.0	0.5	1.0	1.5	2.0	3.0	4.0
	Alpha	0.0	0.3	0.5	0.7	0.9	1.3	1.6
	IR		0.53	0.50	0.47	0.45	0.42	0.39
Strategy	TE	13.6	13.6	13.6	13.6	13.6	13.6	13.6
	Return	2.0	2.0	2.0	2.0	2.0	2.0	2.0
	IR	0.15	0.15	0.15	0.15	0.15	0.15	0.15
Overall	TE	13.6	13.6	13.6	13.7	13.7	13.9	14.2
	Return	2.0	2.3	2.5	2.7	2.9	3.3	3.6
	IR	0.15	0.17	0.18	0.20	0.21	0.23	0.25

Increasing the active risk of the fund has a negligible effect on the overall solvency risk. . However any out-performance feeds directly through to the bottom line. This is due to the effect of diversification.

Of course the above analysis assumes a certain level of positive skill from the fund manager. In practice this will depend on a variety of factors including the quality of the people and the investment process. The trustees (and their advisors) need to take a qualitative view as to the extent to which they believe that active management is rewarded and more specifically in their own ability to judge which managers possess positive skill.

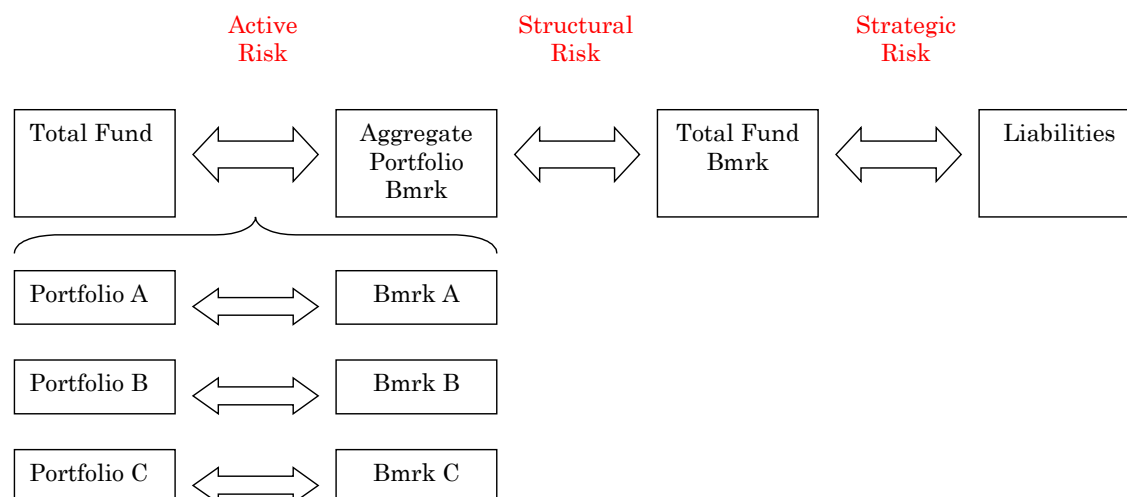
2. Risk budgeting – the consultants’ perspective

2.1 What are the risks?

The conventional approach to risk budgeting (even if it has not been referred to in this manner) with respect to pension funds is scheme-centric. In other words, the pension scheme is regarded as a primary investor and ascribed with attributes such as risk tolerance or utility. The pension scheme, through its agents, the Trustees, must trade-off risk and return; worry about solvency and ongoing contribution rates.²

From this starting premise, the risk budget is then usually decomposed into ‘strategy’ risk and ‘manager’ risk. For the purposes of this paper we think of strategy risk as being the risk associated with possible outcomes from the decisions about which asset classes to include in the benchmark. The active risks at portfolio level are the risks that the active managers take in order to achieve their target levels of benchmark outperformance. In addition there may be some ‘structural risk’ associated with any mismatch between the aggregate of the portfolio benchmarks and the total fund benchmark.

The overall risk (in a scheme centric sense) is the ‘sum’ of the manager, structural and strategic risks. The risk budgeting process in this case is two-fold; first, deciding how to allocate the maximum permitted overall risk to total fund active risk and strategic risk; second, allocating the total fund active risk budget across the component portfolios.



The input that the consultant has varies from scheme to scheme. The trigger for a review of investment arrangements can be one or many of

- an actuarial valuation that has produced an unexpected funding level
- a period of ‘underperformance’ by one or more of the investment managers
- a period of volatility in the investment markets that has hit the media and brought the investment aspects of the pension fund to the trustees immediate attention

² The alternative view is that the investment policy of the scheme should be viewed from the perspective of the shareholders of the sponsoring company. We discuss this view in more detail in section 2.2.

- regular reviews of strategy, structure and managers undertaken by the scheme.

2.1.1 Active risks

Until fairly recently, far more attention (and money) has been lavished on measuring portfolio active risk than on any of the other risks. Although tracking error, warts and all, has been elevated as the queen of risk measures, a vast array of such measures exists: active and common money; downside tracking errors; value at risk; shortfall probabilities; expected shortfalls; etc. A perhaps even more bewildering array of skill measures (risk-adjusted returns) has evolved to enable investors to decide whether or not the active risk generates value.

Empirical evidence would suggest that, for most pension schemes, the total risk (in tracking error form) associated with active management and manager selection is of the order of 0%-5% per annum, with the majority of schemes having risks measured at around 2% p.a..

2.1.2 Fund structuring

For schemes that are not very small, structural risks (in the sense that the portfolio benchmarks do not add up to the total fund benchmark) can be made very small, particularly if peer group benchmarking is avoided for any of the portfolios or for the total fund. Being able to specify the risk is one of the arguments for avoiding peer group benchmarks.

Fund structuring is jointly concerned with the dividing up of the total fund benchmark into portfolios that can be practically managed and the aggregation of portfolio active risks into a total fund active risk. With the ever-increasing range of products offered by the investment management community, the disaggregation of the total fund benchmark is often not a problem for all but the smallest schemes.

One of the areas where a consultant can add value is in finding a combination of portfolio benchmarks that not only adds up to the total fund benchmark, but also does not unduly constrain the active managers' processes. For example, many equity managers have a global research and portfolio construction process and so asking them to manage a portfolio that excludes a particular region, e.g. the UK, will constrain them. The most common approach is to use regional indices. Alternative approaches include the use of indices focusing on global sectors or the use of multinational and local indices.

Although some of the more sophisticated risk attribution systems were originally designed to help fund managers more accurately understand the active risks that they were taking, they also provide valuable information to trustees and their advisers in structuring the fund. The active risk attribution enables 'complementary' managers to be identified. The style of the manager (e.g. growth/value) can be picked out from their risk signature. Style identification can be used to ensure that the best manager within each style is picked and yet keep the fund as a whole style-neutral (unless the trustees particularly want to take on a style risk).

Although in the UK these styles are not always very distinct, it does enable the consultant to check that the total fund is not over-exposed to any one style when the scheme employs many managers.

Consultants will have in-house models, or use third-party systems, to structure the fund in an 'optimal' way within the scheme's total fund active risk.

2.1.3 Strategy

There are various approaches to recommending strategy in practice, but most of them involve

- specifying a set of asset classes (e.g. equities, bonds, property, etc.)
- making assumptions about their 'long-term' investment characteristics, e.g. returns, correlations, variances, etc.
- finding a set of 'efficient' portfolios from these asset classes to be candidate strategies subject to client-imposed (and often adviser-imposed!) constraints
- describing the risk and return characteristics and consequences of the candidate strategies on the fund
- somehow choosing among the candidate strategies

Until recently it has been very difficult to place a value on the strategic risks. This is because the value of the liabilities has been ill-defined. They have been creations of the actuarial profession's judgement and have moved from valuation to valuation in seemingly mysterious ways.

In the last 5 or so years, various developments have meant that the liabilities' value has become a little more identifiable. The MFR was introduced and provided a standard and very prominent way of valuing the liabilities; the markets dared to move away from actuarial theory and a credibility gap opened forcing actuaries to adopt a more market related approach to valuations. The new accounting standard, FRS 17, provides yet another way of measuring liability value.

It was soon noticed that, in comparison with the 2% active risk being taken by the funds, the strategic risks were much greater. The typical tracking error for MFR valuations is between 5% and 10% per annum. Against an FRS17 value of liabilities, the tracking error is more in the range of 10% to 20% p.a. for most schemes.

The actuarial profession (e.g. Urwin et al., 2000) and the regulators (via Mr Myners) have spotted the fact that active risks are typically dwarfed by strategic risks. Indeed, as we have shown earlier, active management risk is virtually negligible in the scheme-centric sense.

In looking at this issue further, we have to unpick what types of risk are being introduced to the scheme at each stage. In setting strategy (policy), the Trustees, often in conjunction with consultants, are mainly deciding on a level of systematic risk (relative to the value of the liabilities).

In reward for taking on systematic risk, the scheme benefits twice. Firstly, the actuarial profession is still able to exercise its judgement in taking advance credit for expected excess performance and thereby specify an improved ongoing funding level and lower ongoing contribution rate. Secondly, over a long period of time, the 'risky' assets do tend to outperform the matching, bond-like assets that (arguably) replicate pension fund liabilities.

The genuine surplus that might therefore emerge (because the actuaries are sometimes conservative as to how much advance credit they do take) can be used to offset the ongoing contribution rate further.

Advisers have typically taken the view that many pension funds should be able to ignore short-term volatility and therefore should be amenable to investing in risky assets, such as equities. The Myners Review makes the case for adopting even ‘riskier’ assets, such as private equity.

2.2 Pension fund investment: latest theory

Finance and economic theory would suggest that, to first order, strategy is irrelevant. The investment returns (and risk) from the pension fund impact on the contributions required to the scheme from the sponsoring company which in turn impact on the company’s profits and hence the return to the shareholders. All that putting systematic risk in the fund does is complicate shareholders’ lives when it comes to establishing their own personal levels of investment risk³.

The pension scheme is not an economic entity and so does not add economic value. It represents just part of shareholder wealth. Shareholders can and do rearrange their assets outside the pension fund in order to establish their level of risk and hence trustees serve the shareholders poorly by agonising over how much or little systematic risk to take on.

If pension funds take on more equity risk to enhance surplus and reduce contribution rates, then shareholders are effectively exposed to equity risk and so must buy more bonds on their personal accounts in order to maintain the same personal level of risk-adjusted return (or utility). Given current taxation regimes, this may represent an economic cost when taken in aggregate.

The return that the pension fund can earn is no more or less than that which the market also offers the individual shareholders. Cumberworth et al. (2000) make a similar point in the context of insurance funds.

Chapman, Gordon and Speed (2001) have demonstrated how differing levels of systematic risk might feed through (mainly because of second order effects) to the various parties involved in a pension fund. Because of taxation differences and member solvency, largely bond-based (as opposed to equity-based) investment strategies have much to recommend them for pension funds (see Exley, Mehta and Smith (1997)).

2.2.1 Where is the added value?

Although this viewpoint would seem to be rather gloomy for consultants and equity managers, it does have a silver lining. Because systematic risk is not fundamentally crucial to the pension fund, active risk comes back into its own.

If active managers can be found who are genuinely adding value by exploiting inefficiencies in the market or removing market frictions and if pension funds have an advantage over

³ See Exley et al (1997), Chapman et al (2001) and Cumberworth et al (2000) for further discussion.

individuals in employing the managers then such investment by the pension fund might be adding economic value. Consequently there is also value in consultants helping pension funds to find such managers.

The advantages that a pension fund might have over an individual in this regard may come from a variety of sources. For example,

- (a) from being able to negotiate lower fees
- (b) because the assets – such as private equity opportunities – are effectively available only to large pools of money
- (c) because institutional investors can more cost-effectively lobby the investment managers to take an active role (or at least interest) in corporate governance
- (d) through greater access to manager research that enables them to identify the more skilful managers.

The “ifs” above are probably quite significant and moreover the economic value added would have to be traded-off against the perceived strategic costs of investing in equity. We make no claim in this paper that this is an overwhelming reason for pension funds to invest solely or even significantly in active managed equity or private equity portfolios, for example.

Since active risks are typically non-systematic risks, shareholders can diversify them. Any ‘alpha’ being generated is then wealth that is added to their pockets.

There is something of an irony here: although active management might well have been supported previously by pension funds for scheme-centric reasons, its continued support may depend on the rejection of the scheme-centric approach and a focus on economic value added.

2.3 Pension fund investment: current practice

The theory outlined above is being hotly debated in actuarial circles. However in practice nearly all schemes remain ‘scheme-centric’ in that they focus on the investment policy of the scheme in an independent manner.

There is an argument that the introduction of FRS 17 may highlight to shareholders the risks that are being run (essentially on their behalf) within the pension scheme. As a result this may hasten the move to a more holistic view of pension fund investment policy.

The alternative view is that pension schemes are relatively comfortable in their current state and the key decision-makers do not have the level of in-depth understanding, time or inclination to change this stance. The impact of behavioural biases in this regard should not be underestimated. In particular it is human nature to:

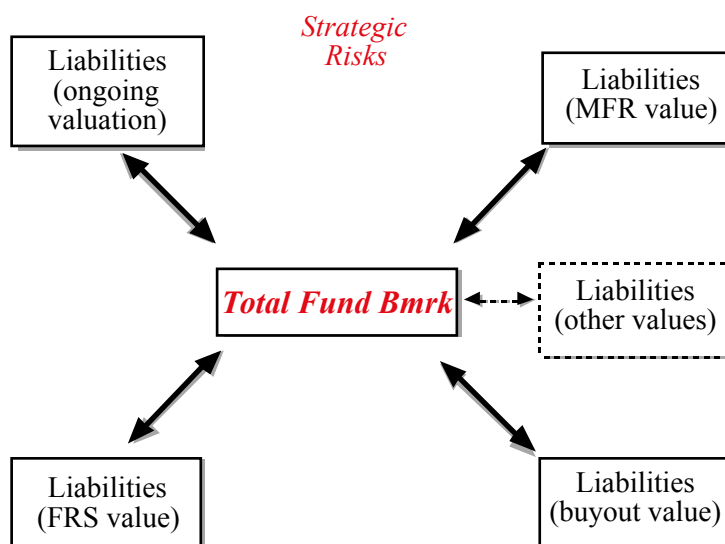
- look to keep things simple by putting the pension fund into a different mental compartment rather than consider the complexities of how shareholders’ interests sum together
- avoid “regret risk” by keeping in line with common practice
- be influenced by recent experience. After all, adopting a mismatched position has (on the surface) served pension funds (and the sponsoring employer) very well over the last 20 years.

There are also concerns as to whether sufficient “matching” assets (i.e. fixed and index-lined bonds) would be available should pension funds decide to go down this route on mass.

2.3.1 Difficulties in defining the risk budget

One of the recommendations of Mr Myners is that closer attention is paid to strategic benchmark setting because that is the big picture in investment strategy. However, by encouraging focus on an activity that arguably has no meaningful economic value there is a chance that the area where economic value may be being added is destroyed.

Focusing on strategy engages the investment consultant in a spiral of second-guessing what the scheme actuary might do. The scheme sponsors pay into the fund a rate that depends on the actuary’s contribution rate that in turn depends on his or her assumptions about investment returns. The trick is to persuade the actuary that the investment strategy will produce a steady stream of high returns so that he can certify a better current funding position and also a lower future service contribution rate.



But with the introduction of other measures of solvency in addition to the ongoing funding position (e.g. MFR, while it still lasts, and FRS17, when it comes into play) makes the ‘optimal’ solution poorly defined, if not non-existent. It is therefore (practically) impossible to specify unambiguously what a good strategy is and what a bad strategy is, except after the event.

It is even more difficult to specify what the ‘risk’ of a strategy is. The risk of the strategy is at best multi-dimensional and is usually dynamic. Risk budgeting is therefore extremely difficult from a scheme centric point of view. The risk budget is multi-dimensional and ‘weights’ have to be put on each of the dimensions to ascertain their relative importance for the scheme.

Although there is little mathematical trouble in ‘spending a multi-dimensional risk budget’, the specification of the budget and the spending process will depend on weights that are difficult to interpret.

Consultants advising on strategic risk are therefore faced with an extremely difficult task since the total overall risk budget is ambiguous. In practice, they often have to resort to 'experimenting' with a lot of different strategies and illustrating their possible consequences on funding levels (using the various liability measures) and recommended contribution rates.

The recommended strategy is then the one where none of the strategic risks is too great and the trustees are happy that they can persuade the actuary to continue recommending a low-ish ongoing contribution rate.

3. Risk Budgeting – A Fund Manager’s Perspective

One of the key aspects of investment strategy for any fund is to decide the levels of risk that it is appropriate for the fund manager to take. This needs to be defined at both total fund level and sub fund level, e.g. UK equities, US equities. This section outlines a framework to assist this decision making process and a method of arriving at an appropriate risk budget for a hypothetical fund, Fund X, which has a benchmark typical of many balanced funds.

There are several aspects to consider:

3.1 Relationship between risk and return

The core assumption underlying the risk budget is the *prospective* relationship between active risk and excess return. The metric of return divided by risk is generally referred to as the information ratio (IR) – it measures the amount of expected return per unit of risk.

Unfortunately, this is a complicated relationship. A sensible starting point is to assume an information ratio (at sub fund level) of 0.3 for small levels of active risk. This may seem small, but the IR will be larger at total fund level because of risk diversification across the asset classes (we will return to this later). In general, this is the observed level of added value for “good” active asset managers, and widely accepted to be an appropriate target.

From within a fund management house, it is sensible to assume the same prospective levels of information ratio for all asset classes whether Japanese equities, property or corporate bonds. To assume otherwise would present some difficult and rather sensitive management issues. If we really believed that we could add more value from UK rather than US equities then it would be logical to hold all assets as UK equities and achieve US equity exposure from short UK and long US futures positions (or some other market instrument, such as a return swap). This is known as “alpha porting” or “alpha transfer”.

In practice, it is very difficult to make such assessments and most trustees are unlikely to be comfortable with the use of alpha porting. It should be noted that an investment consultant, unconstrained by choice of asset manager, might make very different assumptions and arrive at correspondingly different answers.

3.3 Information ratios for differing levels of risk

A complication arises because of the practical constraints imposed on fund managers. If a stock has a 1% index weight, then it is possible to have a long position of +2% but not the equivalent short position. We are constrained to have a minimum holding of nil, and therefore a maximum short position of –1%. In constructing a portfolio, the fund manager will typically have a range of views that can be ranked in order of expected value added. The “best” views will be added to the portfolio until they can no longer be implemented, then the second best

and so on. As more risk is taken, the further down the list the fund manager will have to go. Thus, as risk is increased, the information ratio would be expected to fall.

This general concept is encapsulated in portfolio theory and is more commonly referred to as the “fundamental law of active management”. This is discussed further in Appendix A.

But how does the information ratio decline with levels of risk? One would expect the result to differ according to the shape of the benchmark. The S&P500 would be expected to be less constraining than the 19 stock index for FT/S&P Netherlands, where Royal Dutch comprises over 30% of the index. This is not an area that has been widely researched yet is at the core of the investment process, with many implications for marketing statements. However, we do have some insights from a study by Barrie & Hibbert⁴.

The Barrie & Hibbert approach was to simulate 50 stock active portfolios, making assumptions about the inherent skill of the manager and different benchmark shapes within a 50 stock index. Stock returns were randomly generated and portfolios at different levels of risk were constructed using an optimiser. Although this is an artificial construction, which bears little relation to any particular fund manager portfolio construction processes, the general results are a useful starting point.

Chart 1 : Simulated information ratio profiles for top 50 FTSE weights

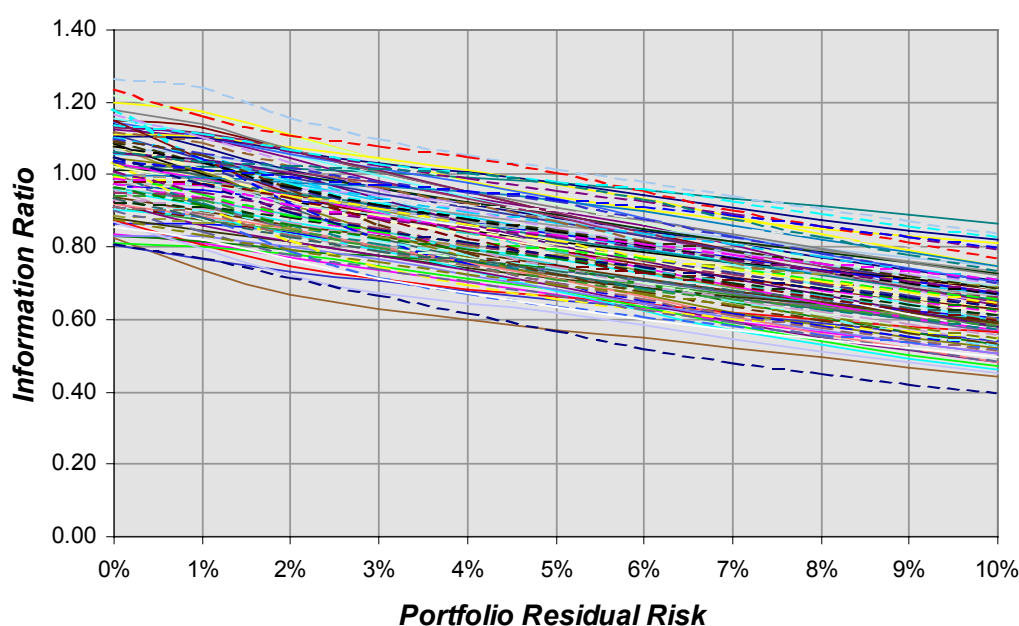


Chart 1, sourced from the Barrie & Hibbert paper, gives a typical set of results for a manager with an underlying information ratio of 1. These results can be scaled in accordance with our own beliefs. Although based on FTSE, the decline in information ratio follows very similar patterns for other benchmark shapes. Each line in the chart represents a different simulation, which produces a spread in observed information ratio around the central assumption. There is a clearly observable declining trend in information ratio as the residual risk (tracking error) increases. Effectively, the short sale constraints act as a drag to achieved performance. On average, the information has declined by around 40% as the tracking error approaches 10%.

⁴ The Efficient Frontier in Stock Portfolios, Barrie & Hibbert, December 1997

Based on such results, the analysis that follows uses a general rule of thumb is that the information ratio will halve by the time the active risk (tracking error) has reached 10%. The assumption of a 50% decline in information ratio gives a margin for prudence over the Barrie & Hibbert results. We will therefore use as our core illustrative assumption an information ratio of 0.3, declining linearly to 0.15 for a 10% tracking error.

3.4 Relationships between sources of added value

Another difficult issue to consider is the relationship between sources of excess returns. This will depend on the approach followed by the investment manager.

For investment managers who incorporate strong “top-down” views in their portfolios we would expect to see some correlation between the excess returns from different teams and between the asset allocation and stock selection process. For example house views on interest rates, say, are likely to be implemented in both asset allocation, bond and equity portfolios.

Further, there may be a house style of investing (such as “value” or “growth”) that has been implemented globally across all equity portfolios. It would therefore seem that the returns across regions would not be independent unless style returns were a purely local phenomenon. Recent history of the rise and fall of TMT stocks across global markets provides some anecdotal evidence that this is not the case.

For investment managers who adopt a “bottom-up” stock-picking approach it is less likely that we would see correlation between the excess returns of the different teams. This is because each team’s portion of the portfolio comprises a number of (largely) independent stock decisions.

The conclusion is that the level of assumed independence of sources of added value will depend on the manager’s investment process. Research would need to be done to confirm the appropriate level for any particular manager. For the purpose of setting the risk budget in this example we initially assume an average correlation of 0.2 across asset classes, and then re-run the problem assuming zero correlation.

3.5 Optimisation

If we take the key assumptions outlined above, then the solution to the risk budget problem can be found by optimisation. Given the set of portfolio weights, we are looking for the set of risks across asset classes that maximises the total expected return, and hence the expected information ratio. A range of solutions are shown in Table 1 below for different levels of total risk using Fund X benchmark weights.

Table 2: Risk Budget for Fund X, for different levels of active risk, assuming 0.2 correlation

	wt%	risk=t.e.	risk=t.e.	risk=t.e.	risk=t.e.	risk=t.e.	risk=t.e.
AA		0.1	0.4	0.6	0.9	1.6	2.3
UK Equity SS	54.2%	0.3	0.6	1.1	1.5	2.4	3.3
Eur Equity SS	10.0%	1.1	2.1	3.0	3.7	4.8	5.7
US Equity SS	4.5%	1.8	3.0	3.9	4.6	5.5	6.2
Jap Equity SS	3.6%	2.0	3.2	4.1	4.7	5.6	6.3
Pac Equity SS	2.6%	2.3	3.5	4.4	4.9	5.8	6.4
other Equity SS	1.2%	3.0	3.7	4.5	5.3	6.0	6.6
UK Bond SS	8.9%	1.2	2.2	3.2	3.8	4.9	5.8
OS Bond SS	4.9%	1.7	2.9	3.8	4.5	5.5	6.2
I-L Bond SS	3.0%	2.2	3.4	4.2	4.8	5.7	6.4
Property SS	2.6%	2.3	3.5	4.4	4.9	5.8	6.5
Cash SS	4.5%	0.0	0.0	0.0	0.0	0.0	0.0
	100.0%						
Total fund t.e.		0.50	1.00	1.50	2.00	3.00	4.00
inf ratio		0.53	0.50	0.47	0.45	0.42	0.39
alpha		0.26	0.50	0.71	0.90	1.26	1.57

The effect of risk diversification is clear. Whereas the maximum assumed information ratio at the asset class level is 0.30, at total fund level this has increased to around 0.50. The actual level declines from 0.53 (for total fund tracking error of 0.5) to 0.39 (for a total fund t.e. of 4.0). Across this range the expected returns (alpha) from active management range from 0.26% pa to 1.57%pa.

In general, the smaller the weight of an asset class within the overall portfolio, the larger the optimal tracking error. This is an intuitive result and bears out what is often referred to as “peripheral concentration” i.e. that we should be more relaxed about taking risk in markets that are a small percentage of the benchmark to avoid undue dilution of potential returns. The optimisation gives us a framework to help us choose some specific parameters.

In practice, it may be difficult to achieve the levels of active risk within the bond portfolios shown in Table 1 without taking significant credit and/or currency risk. For similar reasons, the tracking error of the cash component has been constrained to zero. Outside the UK, equity risks are similar across the different regions with slightly lower numbers for Europe, reflecting the higher weighting. The asset allocation target risks move in line with total tracking error, as one would expect.

In passing, it is interesting to note that if we assumed independence of risk between asset classes, the equivalent numbers for total fund risk would be:

Table 3: Risk Budget for Fund X, for different levels of active risk, assuming zero correlation

	wt%	risk=t.e.	risk=t.e.	risk=t.e.	risk=t.e.	risk=t.e.	risk=t.e.
AA		0.2	0.5	0.9	1.3	2.1	2.9
UK Equity SS	54.2%	0.4	0.9	1.5	2.1	3.3	4.3
Eur Equity SS	10.0%	1.7	3.5	4.9	5.9	7.2	8.1
US Equity SS	4.5%	3.2	5.3	6.8	7.6	8.5	9.0
Jap Equity SS	3.6%	3.7	5.9	7.3	8.0	8.8	9.2
Pac Equity SS	2.6%	4.5	6.8	7.8	8.5	9.2	9.4
other Equity SS	1.2%	6.4	7.7	8.4	9.2	9.5	9.7
UK Bond SS	8.9%	1.9	3.7	5.2	6.2	7.5	8.2
OS Bond SS	4.9%	3.0	5.3	6.5	7.5	8.4	8.9
I-L Bond SS	3.0%	4.1	6.4	7.6	8.3	9.0	9.3
Property SS	2.6%	4.5	6.8	7.8	8.5	9.2	9.4
Cash SS	4.5%	0.0	0.0	0.0	0.0	0.0	0.0
	100.0%						
Total fund t.e.		0.50	1.00	1.50	2.00	3.00	4.00
inf ratio		0.84	0.74	0.66	0.61	0.53	0.48
alpha		0.42	0.74	0.99	1.22	1.60	1.91

The expected returns and information ratios are significantly higher and probably rather less realistic.

It is important to note that the above figures are all gross of fees. The effect of fees will depend on the size of the fund and on the approach adopted. In particular larger funds will tend to have lower fee levels and fees will typically be higher when specialist managers are used rather than using one manager for the whole fund.

In practice we recommend that net of fee figures should be used for this analysis. When assessing the net of fee figures it is important for the trustees to bear in mind the fees associated with the passive alternative given that there is no cost-free option.

3.6 Return Target

The key variable remaining is to decide the overall target risk levels for Fund X. It may be the case that this risk target is merely handed down following the type of process outlined in section 2. However, most current practice suggests it is generally easier to think of this in terms of the return side of the equation, such as setting a target performance of, say, 1% per annum versus the benchmark index. From Tables 2 and 3 above a 1% return target implies a total fund target tracking error of above 1.5% but under 3% for the fund, depending on the views on independence of the alpha generation process across asset classes.

This result is broadly consistent with the general rules of thumb that are in operation. A typical view is that the tracking error target should be around 2 to 3 times the performance target, i.e. based on an overall information ratio of approximately a third to a half. General perceptions are that targeting an information ratio of greater than a half is somewhat

unrealistic (over the long-term) and increases the likelihood of disappointment. At the other extreme targeting an information ratio below a third might not be regarded as active management and any outperformance may struggle to cover the fee.

In passing, note that it is important for trustees to focus in on the information ratio that they believe a manager may realistically achieve rather than what they say they might achieve. For example, manager A and manager B may be pitching against each other for a mandate. Both are proposing to manage the fund with a tracking error around 3%. Manager A states their target outperformance as 1.5% p.a., whereas manager B is targeting 1.0% p.a.. On the surface manager A seems more attractive, but how confident should the Trustees be that manager A has more skill than manager B and will deliver higher prospective returns? . There is no easy answer to this, but the Trustees should be aware of the dangers of being deceived by marketing “spin”.

4. Managing Active Risk

Having set the active risk budget for the fund (or portion of the fund) the trustees then need to ensure that the fund is being managed in line with this budget. In this section we outline our views on what trustees should be looking for from their investment managers in this area.

The portfolio risk management⁵ process of the fund manager should include:

1. Clear, well thought-out guidelines on the risk profile of the fund
2. A clear control process establishing when a fund is outwith its range, the appropriate action to be taken and the responsibilities of all involved.

4.1 Establishing the risk profile

There are a number of measures that can be used in establishing the risk profile of a fund. Typically a target tracking error range will be used as the main focus. This may be supplemented by guidelines on stock, industry, country and style exposures.

If the trustees have followed the risk budgeting process outlined in section 1 of the report then there will be a tracking error target for the fund as a whole and for specific active managers (if more than one). Alternatively, as noted at the end of the previous section the risk target may be inferred from a performance target.

In practice it is more appropriate to give investment managers a broad tracking error range than a specific target. There are a number of reasons for this:

Firstly, as we have seen earlier, the active risk is typically swamped by the strategic risk. The overall solvency of the fund is not particularly sensitive to changes in the active risk level. For example, if the investment manager were to have an extreme level of underperformance in any one year then the impact of this is likely to be swamped by the absolute performance of the strategic asset allocation.

Secondly, a narrow tracking error range can encourage the manager to adopt a myopic focus and an over-zealous approach to risk management. In this scenario risk considerations can drive portfolio decisions and force turnover in order to maintain the target risk level. This will not only incur transaction costs but may also distract the manager from their key goal of adding value.

A further reason that the trustees may choose to have broad tracking error ranges is to allow the investment manager some freedom to increase or reduce the risk level based on the conviction of their views. This increases the scope for the manager to outperform.

⁵ We specifically refer to “portfolio” risk here to distinguish this from other areas of risk facing an investment manager (e.g. operational risk, regulatory risk, etc). Portfolio risk may also be referred to as investment risk and specifically focuses on managing the range of returns of a portfolio. This involves balancing the conflicting goals of minimising the risk of extreme underperformance whilst allowing sufficient scope for the manager to achieve the performance objective.

Finally, the prospective tracking error of the fund is not directly observable or controllable. Any estimate of the future tracking error is subject to a high degree of error. Focusing in on a specific point target is therefore spurious and doomed to failure. An awareness of this fact at outset can save much angst at a later date.

In practice there are a number of psychological factors that encourage funds to be managed within low, narrow tracking error ranges. By allowing broad risk ranges there is an increased risk of extreme short-term underperformance. This can be uncomfortable for all involved:

- Although pension funds invest for the long-term it is human nature to look at any short-term performance information available and be overly concerned by periods of short-term underperformance.
- Investment managers are concerned with managing their own business risk. If they feel that trustees will react badly to periods of short-term underperformance then their self-preservation instincts will lead to index-hugging.

The Myners report highlights these issues and emphasises the need to maintain a long-term perspective. We fully support this view.

In terms of the specifics of establishing tracking error ranges for each asset class there are no hard and fast rules. However we can use the analysis in table 3 from the previous section as a guide. For example with a tracking error range of 2 – 3% for the overall fund this implies a range for UK equities of 2.1% to 3.3%. If we are unsure about the correlations of active risk, combining tables 2 and 3 would give us a tracking error range of 1.5% to 3.3%.

Having established a target tracking error range we then need to consider other aspects of the risk management profile. Guidelines on maximum stock, sector, country and region positions are useful supplements and can help guard against some of the flaws in estimating tracking errors discussed in the next section.

For mandates requiring an explicit style bias (e.g. value/growth) it may also be desirable to have guidelines on style measures. However in practice style is a somewhat subjective concept that is difficult to set parameters around.

4.2 Structure of the portfolio risk management process

Having established the risk profile for a fund we then need to ensure that there is a control process in place in order to monitor whether the fund is being managed in line with this profile. The key aspects to such a process are:

- Buy-in from senior management. It is critical that senior management have an understanding of portfolio risk and have approved a formal portfolio risk policy.
- A specialist portfolio risk management department. This team will be central to the risk management process.
- A clear process for identifying and reporting on whether a fund is being managed in line with its risk profile.
- A clear statement of the actions that need to be taken when a fund is outwith its risk range and the responsibilities for ensuring that these actions are followed through.

In terms of the mechanics of the process there are likely to be two elements:

- On-line controls to stop inappropriate deals going through. For example this could be based on stock/sector/country limits plus client restrictions (e.g. no tobacco). This aims to guard against “rogue fund manager” and accidental error concerns.
- A regular reporting process (e.g. monthly) to ensure that funds are within appropriate risk ranges taking into account all relevant measures.

Ideally portfolio risk management should be integrated in the investment process such that fund managers have an in-depth understanding of portfolio risk and are aware of the risk that they are taking when they are placing deals. The monthly control procedures should therefore be a check on this process rather than the sole form of risk control.

One of the more difficult parts of this process is establishing whether a fund is being managed in line with its risk profile. In the next section we consider how the risk level of the fund may be estimated.

4.3 Assessing the risk level of a fund

Some managers will use a risk model as their sole guide as to the risk level of a fund. There is then a very clear guide as to whether the fund is running the appropriate level of risk. The disadvantage of this approach is that the manager relies completely on the accuracy of the model.

The flaws in risk models have been discussed at length in other papers (most notably the paper from this working party last year entitled “Predicted Tracking Errors: Fact or Fantasy”). The key observations of such studies are that:

- Extreme events are far more common than the models predict (especially in recent years)
- The three key causes of these extremes are:
 - Momentum in stock returns
 - Changing volatility levels
 - Changing factor effects
- Predicted tracking errors are less reliable for more concentrated funds

As a result of these flaws it is increasingly being recognised that over-reliance on any one model is dangerous. Managers are therefore incorporating other tools and analysis in order to provide a wider perspective on risk.

Under this approach, the assessment of the risk level of the fund is still likely to be heavily influenced by the output of a risk model. However the wider perspective enables the portfolio risk specialists to use their experience and judgement to express a view on any bias in the model. These additional tools and analysis are discussed below.

1. Use of another risk model

The prediction of any risk model is heavily dependent on the data history used to estimate the volatilities of (and correlations between) the assets. Using an alternative risk model that is based on a different data history (i.e. shorter or longer) can help to give a different perspective.

2. An analysis of past performance data

The rolling 1-year, 3-year and 5-year tracking errors of the portfolio will provide an additional guide to the risk level of the fund. However if the way in which the fund is being managed has changed then this may not be representative of the prospective risk levels.

The performance history may also identify any momentum in the returns of the portfolio, i.e. by observing runs of positive or negative monthly (or quarterly) relative returns. In addition performance attribution can identify the key contributors to any extreme performance levels and aid further understanding of the risks in the portfolio.

3. Risk Attribution

From the output of a risk model it is possible to identify how each stock (and factor exposure) contributes to the predicted risk level. This can act as a check on the portfolio construction process. In theory the highest contributions to risk should reflect the manager's strongest views. It can also help in identifying whether the risk is evenly spread or is heavily weighted towards a few key exposures.

4. Incorporating a short-term risk measure

The volatility of markets and individual stocks will change over time. In theory short periods of higher volatility should not be of concern to long-term investors whose focus is more on avoiding extreme underperformance on an annual (or longer) horizon rather than over one month or one quarter. However an eye on what is happening in the short-term can act as an early-warning signal of longer-term changes.

This can be achieved by:

- using a trading risk model based on a very short (e.g. 20 to 60 days) history of daily data
- using option-implied volatilities as a guide to market expectations and calibrating risk estimates appropriately
- calculating the tracking error of the fund using a relatively short period of daily data (e.g. 20 to 60 days)

The danger of incorporating such a measure is that it can encourage a short-term focus resulting in increased churn of the portfolio. It is therefore important that this information is used with care.

5. Active/Common money

The risk of a fund relates to two elements:

- the aggressiveness of the fund in terms of the size of the bets taken
- the volatility of (and correlation between) the assets held

Active and common money measures provide a simple indication of how aggressively the fund is being managed. Observing how these measures have changed over time can help to identify any narrowing or widening of the stock bets being taken by the fund manager. This is particularly useful in identifying whether any change in the predicted tracking error level is due to the fund manager being more aggressive or due to changes in (assumed) volatility levels. An example of the use of this analysis is attached in appendix B.

6. Style analysis

Style analysis has a number of applications. The main use is as a confirmation of a manager's investment style (e.g. growth/value, small/large). In the US investment managers typically have well defined styles and pension funds will aim to combine managers of different styles in order to maintain a neutral style exposure at the total fund level. In Europe the use of styles is not as developed but with the increasing globalisation of the fund management industry this is likely to change in future years.

Where managers have been selected for a mandate based on a specific style then it is important for the manager to monitor that this style is being maintained. In times when the style is doing poorly there is a temptation for the manager to neutralise the style bet thus giving the overall fund an unintended style bias.

There is also an increasing acknowledgement that style analysis can help in detecting risks within the portfolio that are not picked up by risk models, in particular in relation to momentum.

Although there is general agreement that momentum exists within markets from time to time its occurrence, severity and cause are difficult to predict. A number of practitioners suggest that momentum is style based (see Fishwick). Understanding the exposure of the fund to style biases and the potential impacts of these biases should help in understanding the potential impact of momentum on the portfolio.

An example of the use of style analysis is attached in appendix B.

Analysis over time

For all of the above forms of analysis it is useful to observe how the relevant measures change over time. This can help to identify the stability within a manager's process and the sensitivity of a model or measure to recent volatility levels.

5. Conclusions

Interaction between strategic and active risk

We have demonstrated that for the majority of pension funds strategic risk swamps active risk. This provides support for two of the recommendations from the government-sponsored Myners report:

- “The attention devoted to asset allocation decisions should fully reflect the contribution they can make to achieving the fund’s investment objective.”
- “Where they believe active management to have the potential to achieve higher returns, funds should set both targets and risk controls which reflect this, allowing sufficient freedom for genuinely active management to occur.”

A significant complication is that, given the wide range of solvency measures that exist (e.g. on-going, buy-out, MFR/FRS17), the level of strategic risk is multi-dimensional. This makes it difficult to set (and confusing to understand) the level of strategic risk of a fund.

The financial economics perspective

Finance and economic theory would suggest that, to first order, strategy is irrelevant. The pension scheme is not an economic entity and so does not add economic value. It represents just part of shareholder wealth. A change in the risk and return profile of the pension scheme impacts on the contributions required to the scheme from the sponsoring company which in turn impacts on the risk and return profile of the company’s profits. Although returns to shareholders may increase, these are accompanied by an increase in systematic risk.

Shareholders can and do rearrange their assets outside the pension fund in order to establish their own chosen levels of risk and hence trustees serve the shareholders poorly by agonising over how much or little systematic risk to take on. All that putting systematic risk in the fund does is complicate shareholders’ lives when it comes to establishing their own personal levels of investment risk.

Although this viewpoint would seem to be rather gloomy for consultants and equity managers, it does have a silver lining. Because systematic risk is not fundamentally crucial to the pension fund, active risk comes back into its own.

If active managers can be found who are genuinely adding value by exploiting inefficiencies in the market or removing market frictions and if pension funds have an advantage over individuals in employing the managers then such investment by the pension fund might be adding economic value. Consequently there is also value in consultants helping pension funds to find such managers.

Even if the stakeholders of the pension scheme accept the above theory they may still prefer to take on systematic risk within the pension scheme in exchange for lower contribution rates. Although the financial economics perspective is now becoming mainstream within the

actuarial profession, it remains to be seen whether the wider acceptance of the theory will result in wholesale changes in pension fund investment practice.

Allocating the risk budget within asset classes

Having set an overall active risk budget we have outlined a framework to consider the appropriate levels of risk to run within each asset class. The relationship between risk and return has been explored and a reasonable set of assumptions outlined from the perspective of a fund manager. In particular we have considered the assumptions to be made regarding levels of skill assumed across asset classes and how skill declines as the risk level is increased. Using such assumptions we have demonstrated how to set a risk budget for asset classes using an optimiser, seeking to maximise returns (or information ratio) for a given level of total active risk.

Risk Management and Monitoring

Pension funds should establish risk objectives for their managers and ensure that managers have appropriate procedures in place in order to manage to these objectives. However we advise against an obsessive focus on short-term active risk which may hinder the overall objective of enhancing long-term returns. Given the minimal impact that active risk has on the overall risk of the fund and the limitations in predicting future risk levels the risk objectives should be fairly broadly defined.

A sound portfolio risk management process will look at a range of measures. Whilst recognising the limitations of the tracking error measure we suggest that the fund uses broad tracking error ranges as the primary guide to the level of risk to be run by their manager(s). This should be supplemented by other measures (e.g. guidelines on maximum country, sector and stock positions) as required.

Appendix A

*The fundamental law of active management*⁶

$$IR = IC * \sqrt{BR}$$

where

- IR is the information ratio of the portfolio
- IC is the information coefficient of the manager. This is the purest measure of skill and is equal to the correlation between the securities' actual returns and the manager's forecasted returns on the securities
- BR is the strategy's breadth defined as the number of independent bets taken over a given time period. (A lack of complete independence between the bets waters down the result but the principle remains.)

In practice both the information coefficient and breadth are probably unobservable but it is the principle that is important. As the number of bets decreases (i.e. the portfolio becomes more concentrated) the information ratio decreases. (Note that the definition of breadth includes both concurrent and consecutive decisions.)

⁶ See "Active Portfolio Management" by Grinold and Khan (1995) for further detail

Appendix B

Use of active money measure


The graphs below show the change in predicted tracking error and the change in active money (for the 10 largest overweight positions and 5 largest underweight positions). The predicted tracking error has doubled over the last 7 years. Looking solely at this measure one would infer that the manager is being much more aggressive.

However focusing on the largest 10 overweight positions in the active money chart we can see that the aggressiveness of the manager did increase slightly from 1993 to 1998 but has dropped significantly since then and is clearly lower at the end of the period than at the start⁷.


This apparent anomaly is due to the fact that volatility levels (and in particular stock-specific volatility levels) have increased dramatically in the late 1990's. Broadly speaking a portfolio with the same level of aggressiveness (or bet sizes) is twice as risky now as it was 5 years ago.

This helps to illustrate that by splitting out the risk of a fund into aggressiveness and volatility one can develop a greater understanding of how the fund is being managed.

Predicted tracking error using a risk model




Active money: sum of top 10 positive, top 5 negative positions



⁷ Some will point out that this analysis only focuses on the top 10 positions in the portfolio and may be misleading. More robust measures (e.g. active money squared) provide the same conclusions but we have focused on the top 10 positions for ease of explanation.

Example of Style analysis

The chart below shows an example of style analysis. This is based on a snapshot of a fund relative to the benchmark at a certain point in time. The blue bars on the left are value measures, the green bars in the centre are growth measures and the red, black and yellow bars on the right provide information on other measures such as size and momentum. This fund shows a clear growth bias.



Source: Style Research

References

Barrie and Hibbert, 'A Framework for Portfolio Risk Management and Control', Research article 1998

Barrie and Hibbert, 'The Efficient Frontier in Stock Portfolios', Research article 1997

Chapman, Gordon, and Speed, 'Pension Funding and Risk', Sessional Paper presented to Institute of Actuaries, April 2001

Cumberworth, Hitchcox, McConnell & Smith (2000), 'Corporate Decisions in General Insurance: Beyond the Frontier', British Actuarial Journal Volume no 6, Part II, No 27

De Bever, Kozun, Viola, Zvan (2000), 'Pension Risk Budgeting: Something Old, Something New, Something Borrowed...', Journal of Performance Measurement, Summer 2000

Exley, Mehta, and Smith, 'The Financial Theory of Defined Benefit Pension Schemes', BAJ, Vol. 3, 1997

Gardner, Bowie, Brooks and Cumberworth, 'Predicted Tracking Errors: Fact or Fantasy', Portfolio Risk and Performance Working Party, Faculty and Institute of Actuaries Investment Conference 2000

Grinold and Kahn, 'Active Portfolio Management', Richard D Irwin Inc. publishing 1995

Urwin, Breban, Hodgson and Hunt (2001), 'Risk Budgeting in pension investment', Actuarial working party 2001