

RISK BUDGETING IN PENSION INVESTMENT

BY R. C. URWIN, S. J. BREBAN, T. M. HODGSON AND A. HUNT

[Presented to the Faculty of Actuaries, 19 February 2001]

ABSTRACT

This paper extends the concept of investment efficiency from investment management structures to include strategic asset allocation and liability related issues. The concept of risk budgeting is developed. It represents a valuable way of incorporating risk and return information to produce more efficient investment decisions. Information ratio is a key measurement in the process, and it is concluded that the risk budget should be allocated based upon the marginal contribution to it for different sources of risk. Non-financial risk is also considered in terms of both governance and risk.

KEYWORDS

Active Return; Governance; Implementation Returns; Net Information Ratio; Regret Risk; Risk Budgeting

CONTACT ADDRESS

R. C. Urwin, M.A., M.Sc., F.I.A., Watson House, London Road, Reigate, Surrey RH2 9PQ, U.K. Tel. +44(0)1737-241-144; E-mail: roger.urwin@eu.watsonwyatt.com

1. INTRODUCTION

1.1 The ‘investment efficiency’ of institutional funds was defined by Hodgson *et al.* (2000). That paper considered the success factors of managing institutional investment funds, and discussed three elements contributing to success:

- the best financial measures of success (principally the *net information ratio*);
- non-financial influences on decisions and results (principally the influence of *regret risk*); and
- the linking factor of governance, and how good governance increases the likelihood of financial success.

The paper considered these concepts specifically in the context of *investment management structures* — the allocations of assets and mandates to different types of investment managers. This paper extends the concepts to asset

classes, with adjustment for strategic asset allocation and liability-related issues in defined benefit pension funds.

1.2 In Section 2 we discuss pension fund governance, introduce the concept of the *governance budget*, and consider how this can be improved. Previous work has identified the considerable influence that a fund's governance has upon investment efficiency. Governance is uniquely defined for each pension fund, and is a key factor in the likelihood of performance expectations being met. Therefore, we start our treatment of this investment topic with a discussion of the problems of organisational design.

1.3 We go on to look at financial measures of success in Section 3. We introduce the concept of *risk budgeting* applied to the pension fund investment programme, and explain its role in the asset planning cycle. Risk budgeting is the assessment of the amount of risk to be employed, and where it is applied. The use of the net information ratio (*active return* divided by *tracking error*) for investment management structures has an analogue in the analysis for the whole fund, in the use of the net *asset liability return* (A/L return) divided by the *tracking error of the asset liability return* (A/L risk).

1.4 Section 4 considers the setting of asset allocation policy in the context of the governance and risk budgets. We assess the *policy returns* arising from strategic asset allocation alongside the policy risks. Section 5 revisits manager structures and considers *implementation returns*. In Section 6 we discuss the practical aspects of risk budgeting, covering the challenging areas of setting assumptions. We consider the vexed question of whether funds should concentrate more of their risk budgets on achieving better policy returns or implementation returns. Our conclusions are found in Section 7.

2. GOVERNANCE

2.1 *The Pension Fund as a 'Business'*

2.1.1 In the vast majority of cases, the running of a pension fund represents a non-core, but financially important, 'business' of the employer, which, for legal reasons, must use separate (trustee) governance. The consequence of this is that the oversight, or *governance*, of the fund raises several challenges.

2.1.2 Trustees face a number of difficulties in managing pension funds. The pension fund balance sheet is generally substantial in monetary terms, complex and very long term. Trustees must manage this balance sheet in the context of multiple stakeholders and within the constraints of legal and financial regulation. Furthermore, trustees' accountability is blurred by their responsibilities both to beneficiaries and to the employer, and their time is limited by other commitments. There are, therefore, varied and conflicting demands placed on trustees.

2.1.3 As a result, pension funds typically represent a non-core business,

and, in general, trustee boards do not draw adequate resources from their sponsoring employer. To compound these difficulties, the employer's financial interest in the pension fund and its risks are often difficult to define and measure.

2.2 *The Governance Budget*

2.2.1 We define the *governance budget* as the capability of the governance structure employed by a fund, and how efficiently it is operating. We can devise metrics for comparing governance budgets between funds, and also for what is necessary to do an efficient job.

2.2.2 The key elements of measuring the governance budget are the size of resources, the skill level of these resources, and the structure of deploying these resources. These can be summarised as:

- time;
- expertise; and
- organisational effectiveness.

2.2.3 Governance can be increased, primarily, by adding resources and by organisation design changes. The application of greater expertise, while naturally desirable, is less easily controlled. The need for additional skill suggests the need to raise the bar in terms of the calibre and the focus afforded to the pension business of a sponsoring company.

2.2.4 Like any managerial resource, the governance of pension funds is limited, but is exacerbated by the pension fund's status as a non-core activity. While internal resources can be supplemented by pools of external resources, without good internal governance these external resources will often be subject to principal-agent conflicts.

2.3 *Good Pension Fund Governance is Difficult to Achieve*

2.3.1 The principles of organisational effectiveness suggest the need for:

- a clear mission or guiding purpose;
- clear and non-overlapping responsibilities and accountabilities;
- separate roles for governing (policy), executive (planning and strategy), and operations (implementation) activities;
- sufficient skilled resources;
- delegation of decisions to those most well-informed; and
- transparency of actions.

2.3.2 Pension fund governance has had six common problems:

- The fund has not operated with a clear mission.
- Trustee boards have undertaken executive roles alongside governing roles, without a clear distinction between the two.
- Decisions have been consensus-oriented rather than given to those best informed.

- External resources have suffered from principal-agent conflicts.
- There have been insufficient resources (both in terms of time and training) to deal with the complexities of pension fund investment.
- The division of time between operational and higher order activities has been skewed to over-emphasise activities that address only the minority of contributions to risk and return.

2.4 *A New Model for Pension Fund Governance*

2.4.1 Drawing on best practice principles from the management discipline of organisational effectiveness leads us to a more detailed governance model, as shown in Figure 1.

- The fund ‘governors’ set the mission and governance arrangements, determine the size of the risk budget, and retain responsibility for monitoring at the macro level.
- The fund ‘executive’ produces the detailed risk budget, and, in so doing, determines the strategic asset allocation, benchmark design and manager structure, and appoints and monitors the investment managers.
- The implementation of the whole programme, that is the day-to-day fund management, custody and performance measurement, is then delegated to fund ‘operatives’.

2.4.2 The new governance model involves clear responsibilities, with lines of accountability from the executive to the governors and from the




Figure 1. New model for pension fund governance

operatives to the executive. However, some level of overlap in the responsibilities is unavoidable, and in several cases is desirable.

2.5 Increasing the Governance Budget

2.5.1 In this model the trustee board operates as the governors; the investment managers and custodians are the operatives; and a specialist asset planning executive must be appointed. The trustee board and the executive should aim to limit any overlap in role or function. The executive needs to be separately resourced, either internally, with an appropriate individual or team, and/or externally, by delegation to asset planning specialists.

2.5.2 This model requires more separation and more specialist resources in the executive function, but, given current overlaps and duplications, this may not produce any additional cost. The model does imply a greater level of management skill applied to the operation of the pension fund business, and, therefore, does produce a considerably higher governance budget than we generally find in the industry as it currently stands.

3. FINANCIAL MEASURES OF ‘SUCCESS’

3.1 Financial Mission

3.1.1 In this section we consider what determines financial ‘success’. Clearly the mission of the pension fund should guide the measures of success. A good mission statement will set down:

- key financial goals;
- secondary financial goals; and
- measures of risk.

These will be appropriate to the fund and well founded.

3.1.2 Success in pension fund investment involves meeting defined liabilities, so measures of financial success must incorporate the assets and liabilities of the balance sheet. All else being equal, success, therefore, involves strengthening the balance sheet by improving the asset/liability ratio. Historically, for most pension fund environments, the asset/liability ratio has, unfortunately, been subjective and not uniquely defined. Furthermore, this objective is not entirely universal as, for instance, too high a ratio of assets to liabilities could imply an unequal treatment of successive generations of members. In this paper we do not attempt to address directly this issue of embedded options in pension fund finance reflecting the conflicting interests of different sets of stakeholders; instead, we investigate the normal situation of moderately funded pension funds.

3.1.3 In this section we formulate a view as to how the key measure of success can be defined, how risk should be budgeted between the key drivers

of success, and consider the factors controlling how large the risk budget should be for a pension fund.

3.2 *The Investment Decision-Making Process*

3.2.1 To fulfil the pension fund's financial mission, there are two possible approaches, which we term traditional and non-traditional.

3.2.2 The traditional, or benchmark, approach splits the investment decision into two parts, the first being the setting of a *strategic asset allocation* or *policy*, and the second being investment management *implementation*. This is the most obvious way to manage a fund with equities and bonds. It involves the creation of a *benchmark* to represent the policy, establishes an investment neutral investment portfolio, and allows the performance from the investment management implementation to be benchmarked against a fair index. The fund's result comes in two parts: the benchmark return; plus the implementation relative return.

3.2.3 The non-traditional, or absolute return, approach involves no benchmark, merely one or more performance targets and risk targets. This is the more natural way to manage a fund investing in alternative assets such as private equities and hedge funds. In such cases there are no obvious benchmarks, only performance comparators. (In our work we find the distinction valuable between a *benchmark*, which is an investable portfolio, and a *comparator*, which is not investable.)

3.2.4 In our methodology, we propose the use of a hybrid method which is driven by the asset classes that are used in the strategy.

3.3 *The Risk/Return Trade Off*

3.3.1 All else being equal, higher returns come with higher risk. The financial mission must consider what balance of risk and return is desirable, taking into account all stakeholders' interests.

3.3.2 The first question to arise is what is considered to be risk. In the context of a pension fund, and given a particular asset/liability ratio, a useful measure of risk is the expected variability (standard deviation) of the future asset/liability ratio. In this paper we only consider risks arising from investment outcomes; risks arising from demographic, legislative and non-investment related factors are excluded from our analysis, as these are largely outside the control of the trustees. For the majority of pension funds, these non-investment related risks are considered 'inherent', and, therefore, must be reacted to rather than planned for.

3.3.3 The risk measure should recognise the log-normal characteristics of the asset/liability ratio which corrects most of its skewness. Semi-variance measures could be considered, but the additional statistical complication does not generally yield much additional benefit.

3.3.4 Other downside measures over alternative time horizons also do not generally add any new information (although they may add perspective

to risk). Essentially, one unit of risk stays as one unit of risk whatever the time-scale, if mean-reversion in markets is assumed to be of minor consequence.

3.3.5 Considering these points, the two key measures are:

- returns; total returns relative to liabilities, or *A/L returns*; and
- risks; standard deviation of [log]return relative to liabilities, or *A/L risks*.

We should also calculate the information ratio (IR), which is the ratio of the A/L return to the A/L risk. This gives a combined measure of the financial efficiency of a policy by assessing the performance per unit of risk. IRs can be compared to gauge financial efficiency at any defined level of risk.

3.4 *The Risk Budget*

Given these metrics, two questions present themselves:

- the first concerns *how* much risk should be taken; and
- the second concerns *where* the risks should be taken to produce better returns.

Together these aspects constitute the *risk budget*, and describe how it is used.

3.5 *The Risk Budget for United Kingdom Pension Funds*

3.5.1 Risk budgets must be assessed relative to a measure for liabilities. For U.K. pension funds there are various measures of balance sheet strength, including ongoing funding, discontinuance funding, the Minimum Funding Requirement (MFR), and accounting (FRS 17). Those like the MFR and FRS 17 have the advantage of being largely objectively defined, and thus capable of more precise analysis. By contrast, many traditional ongoing funding methodologies present problems of estimation, particularly those which are essentially ‘off-market’, and are related to methods based on equity dividend models.

3.5.2 Amongst the various measures listed above, the new accounting standard FRS 17 has particular merits. It is a prescribed and uniformly applicable basis, and thus can be modelled satisfactorily. It also represents the key interest of the sponsoring employers, which, in balance of cost pension funds, have the over-riding concern with risk. In addition, it aligns reasonably well with the revised MFR (as proposed) and aligns to some degree with measures that would be calculated on a gilts plus/market value funding basis. So, while risk budgets can naturally be assessed on any asset liability measure, we see particular merits in using FRS 17 as a reliable basis.

3.6 How Large should the Risk Budget be?

3.6.1 The amount of risk capital that the sponsoring employer and the trustees are prepared to take on is intimately tied to the mission behind taking on this risk. Generally, this is to produce longer-term investment returns, which subsequently lower pension costs and raise benefits per unit of cost.

3.6.2 There are two critical points about the risk budget in assessing how large it should be:

- The first is that the quantum of return that can be produced from the risk budget needs to be considered.
- The second relates to the fact that taking any additional risk in a pension fund must be carefully justified, as the fund is not usually regarded as a ‘core’ activity.

3.6.3 The issue of how much risk is to be taken can be assessed by the stochastic modelling of future fund outcomes in an asset/liability study. These stochastic outputs can be used to quantify the likelihood and size of potential cash consequences for a sponsoring employer, and thereby deduce acceptable and unacceptable outcomes.

3.6.4 There are four fund-specific factors that might influence the individual appetite for risk taking and support different risk levels.

Table 1. Fund-specific factors and influences

	Factor	Influence
Covenant		
	the employer/sponsor covenant to meet future funding	the stronger the covenant, the more risk can be taken
Maturity		
	the term of the liabilities and the period of future funding	the longer the funding period, the more risk can be taken without compromising the security of final benefit payments
Surplus		
	the current funding cushion: current assets minus liabilities	the larger the funding excess, the more risk can be taken
Risk beliefs		
	the subjective view that the trustees have about risk and return	the stronger the risk preferences and/or return beliefs, the more risk can be taken

3.6.5 Each of these factors should be reviewed with the various pension fund stakeholders to reach agreement on the size of the risk budget, and how this should be adjusted over time and in different situations.

3.6.6 Once the size of the risk budget has been determined (a governing board decision), the question of where it is spent can be decided (an executive decision). The shorthand for the two major areas where risk can be taken are policy and implementation:

- *strategic asset allocation and benchmark design or policy*; a matrix of asset classes, percentage allocations and corresponding benchmark indices; and

- *investment manager structure and manager selection or implementation*; a matrix of investment managers or manager types, percentage allocations to each, and performance benchmarks/targets/controls for each investment mandate.

3.7 *Non-Financial Factors and the 'Theta' Budget*

3.7.1 Trustees require other payoffs that are non-financial. Funds tend to give up financial efficiency for non-financial payoffs. See Hodgson *et al.* (2000) for descriptions of the types of non-financial payoffs which we term *theta factors*.

3.7.2 The principal non-financial payoff is control over regret risk (the SleepWell factor), where regret risk is defined as the risk of taking actions that differ from an accepted norm or baseline position. Often this regret risk is assessed relative to the position of not changing strategy, or relative to the industry average position. Examples include the risk of adding to foreign equity exposure, which normally involves adding regret risk, but generally reduces A/L risk.

3.7.3 Such factors are material distortions to financial efficiency to be recognised in any decision framework. In general, funds should seek to limit SleepWell considerations to the minimum necessary for fiduciary purposes. It is unrealistic, however, to assume that SleepWell factors can be eliminated. The consideration of such non-financial influences is, therefore, necessary in our model.

4. DETERMINING STRATEGIC ASSET ALLOCATION

4.1 *'Optimal' Strategic Asset Allocation*

4.1.1 As described earlier, there are several measures of 'return' and 'risk' for a pension fund. Efficient strategies should attempt to maximise future A/L returns at a given level of A/L risk.

4.1.2 There are two choices for fixing an optimal policy. The first is to optimise relative to a principal measure, and then check whether the results are satisfactory relative to secondary measures. An alternative option is to use a 'balanced score-card' to weight the attributes of various policies relative to multiple measures. Stochastic asset/liability methods will be central to this process in either case.

4.2 *The Hierarchy of Decisions in Strategic Asset Allocation*

4.2.1 A certain hierarchy of decisions is needed to produce a detailed policy benchmark:

- The first hierarchy level is the decision of the split between equities and bonds.

- The second hierarchy level relates to the subdivisions of equities and bonds. For equities the subdivisions include geographical and currency splits, and any other type of split, such as size, style or sector. For bonds, the subdivisions include the split between fixed and index-linked, the geographical and currency split, the duration of the benchmark and the credit quality permitted.
- The third hierarchy level concerns allocations to alternative assets, primarily private equity and hedge funds.

4.2.2 The first of these three decision levels is generally the most useful in targeting the overall level of risk to the desired risk budget. The second and third levels of decision can then be used to try to maximise the IR, subject to the desired overall risk level. Generally, the first and third level decisions are viewed as strategic asset allocation, while the second level decisions are viewed as benchmark design.

4.2.3 This hierarchy is consistent with a three-part structure of strategic asset allocation, as shown in Figure 2.

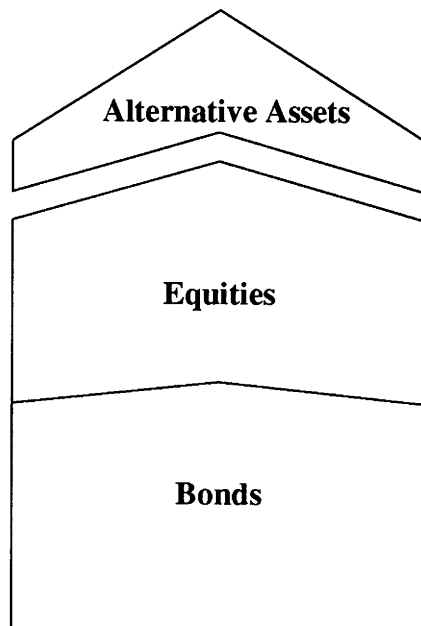


Figure 2. The strategic asset allocation

4.3 *Types of Asset Liability Assets*

4.3.1 The strategic asset allocation in this model is divided into three broad asset classes. Each asset class has a specific role or purpose.

4.3.2 *Bonds*

The role of bonds is to provide safety to the fund. To do this the bonds need to broadly match the liabilities. Government bonds (in appropriate fixed and index-linked proportions) are the baseline strategy to establish the policy risk budget. In benchmark design, there are various opportunities to improve IRs. These include:

- the substitution of non-sovereign bonds (loosely termed credit); and
- the inclusion of foreign bonds with currency hedging.

4.3.3 *Equities*

The role of equities is to provide returns over bonds. To do this, equities would be expected to include U.K. and foreign stocks in widely diversified portfolios. In benchmark design, there are also a number of opportunities to improve financial efficiency and increase IRs, including:

- greater global diversification;
- inclusion of currency hedging; and
- increased diversification through regional or sector controls.

4.3.4 *Alternative assets*

The role of alternative assets is to provide returns above equities and/or risks below equities. The three principal asset classes that provide this mix of attributes are private equity, hedge funds and real estate. The additional return attribute applies most to private equity, the reduction in risk attribute applies most to hedge funds and real estate. The benchmark design opportunities lie in determining the mix of these assets and in the strategic disposition of these assets.

4.4 *The Calculation of Risk Budgets*

4.4.1 The risk budgeting process is central to identifying a policy that carries an appropriate amount of risk and efficient use of that risk. In the following five examples, we illustrate how the hierarchy of policy decisions can be taken in a series of steps.

4.4.2 The calculations are based on the equations set out in ¶3.3.5. The key figures presented in a risk budget are as follows:

- *A/L return* (which we describe as expected out performance); the return premium over the matched or nil risk policy. The calculation is derived from the typical array of assumptions included in an asset/liability study. The calculation we use in the examples below is the annual arithmetic average of stochastic simulations.
- *A/L risk* (which we describe as expected tracking error). The return differences expected relative to the matched or nil risk policy leads to the

tracking error through calculation of the annualised standard deviation of stochastic simulations.

- *Information ratio*. The financial efficiency of any policy (given a certain level of tracking error) is assessed by calculating the A/L return per unit of A/L risk, the higher the IR (after allowing for any costs) the more financially attractive the policy.
- *Expected gain*. The expected out performance can be assessed in monetary terms related to the size of the fund in question. The financial impact of different policies often has most relevance to the trustees and employer.
- *Possible loss* (or VaR). Corresponding to the expected tracking error is the possible loss, which can occur in monetary terms. For risk management and measurement purposes, the VaR is generally defined as the shortfall occurring under a 95% probability, i.e. the monetary value of loss that is exceeded in only 5% of stochastic simulations.

4.4.3 We present the examples of risk budget results for various policy alternatives in Sections 4.4.4 - 4.4.6, 5.2.3 and 5.2.4.

4.4.4 Example 1: risk budget; equity/bond split

4.4.4.1 In this example, we try to understand how the total risk budget is affected by the split between equities and bonds, by considering various allocations. The measurement framework must be chosen; in this example we use FRS 17. This is a prescriptive basis, and implies the nil risk position to be a AA corporate bond benchmark of very long duration.

4.4.4.2 The key economic assumptions required are the asset class returns, their volatilities, and their correlations. We also need the value at risk (VaR) specification, which we take to be the 95th percentile outcome over a one-year period. Furthermore, in all the results shown we assume that the bonds held are long gilts (split 50% index-linked, 50% fixed), and the equities are a typical unhedged 70% U.K., 30% foreign mix.

4.4.4.3 Illustrative results for 50%, 70%, 90% equities policies would be as in Table 2 and Figure 3.

Table 2. Example 1: equity/bond split

Relative to FRS 17	50% equities	70% equities	90% equities
Expected returns			
Out performance (% p.a.)	1.1	2.0	2.9
Tracking error (% p.a.)	10.5	13.6	16.9
Information ratio (%)	10.8	14.9	17.3
Values (£m p.a.)			
Expected gain	11.4	20.2	29.1
Possible loss (VaR)	173.2	224.4	278.2

Note: Fund value = £1bn.




Figure 3. Example 1: risk and return trade-off

4.4.5 Example 2: risk budget; equity/bond sub-division

4.4.5.1 The second example introduces the sub division within the equity component, given the desire to target a risk budget (i.e. VaR) similar to the 70% equity policy in Example 1. This produces more detailed specification of the equity benchmark.

4.4.5.2 In this example, we take the Example 1 70% equity/30% bond split, and test the baseline level of foreign equity exposure, 30%, compared to the alternative strategy of a higher allocation of 50%. In other aspects, we adopt similar assumptions to the previous example. Also, as before, the measurement is primarily with reference to the fund's FRS 17 benchmark. In addition, however, we offer a second measure of risk and return relative to the baseline strategy.

Table 3. Example 2: equity/bond sub division

Relative to FRS 17	Baseline strategy	Alternative strategy	Alternative vs baseline strategy
Expected returns			
Out performance (% p.a.)	2.0	2.1	0.0
Tracking error (% p.a.)	13.6	13.2	1.6
Information ratio (%)	14.9	15.6	2.3
Values (£m p.a.)			
Expected gain	20.2	20.6	0.4
Possible loss (VaR)	224.4	218.6	27.0

Note: Baseline strategy is 21% foreign (i.e. 30% of 70% equities). Alternative strategy is 35% foreign.




Figure 4. Example 2: risk and return trade-off

4.4.5.3 The improvement is more apparent in the main comparison. In relative measurement space, we have to make a judgement as to whether the trade-off is worthwhile, i.e. is the increase in regret risk justified by the reduction in A/L risk?

4.4.6 Example 3: risk budget; allocation to alternative assets

4.4.6.1 The third example considers replacing some of the equities with alternative assets.

4.4.6.2 In this example, we consider a modest level of private equity exposure: 5% relative to the baseline strategy of nil. Again, we adopt similar assumptions to the previous example, the primary measurement is calculated by reference to FRS 17, and the second measure is its position relative to the baseline strategy.

4.4.6.3 Once more the improvement is apparent in the main comparison, while, in relative measurement space, a judgement is required as to whether the trade-off is worthwhile, when considering the increase in regret risk despite the fall in A/L risk.

Table 4. Example 3: allocation to alternative assets

Relative to FRS 17	Baseline strategy	Alternative strategy	Alternative vs baseline strategy
Expected returns			
Out performance (% p.a.)	2.0	2.3	0.2
Tracking error (% p.a.)	13.6	13.4	1.7
Information ratio (%)	14.9	16.9	14.4
Values (£m p.a.)			
Expected gain	20.2	22.6	2.4
Possible loss (VaR)	224.4	221.5	27.4

Note: Alternative strategy is 5% private equity.




Figure 5. Example 3: risk and return trade-off

5. DETERMINING INVESTMENT MANAGER STRUCTURES

5.1 'Optimal' Investment Manager Structures

5.1.1 Given that manager structure is applied relative to a benchmark position, relative performance is the usual measure of 'return', in order to assess optimal manager structures. Risk, here, is best measured by the tracking error (the volatility of the relative return). Efficient manager structure (implementation), therefore, attempts to maximise future returns at given levels of tracking error risk relative to their asset allocation (policy) benchmark.

5.1.2 However, the effect of manager structure on A/L risk is different from its effect on risk relative to an asset class benchmark. Strictly, therefore, we should consider both policy and implementation risk budget figures concurrently, and in the context of A/L risks.

5.1.3 In strategic asset allocation, the choice of separate asset classes is the central issue; in manager structure, the choice of manager types is the equivalent issue. For both, we can consider a certain hierarchy of decisions that produces, successively, a more detailed specification of the fund's structure.

5.1.4 For manager structure the stages are:

- first stage; allocation to passive core;
- second stage; allocations of different asset classes to each manager type: passive, active core and active satellite; and
- third stage; allocations to alternative assets/absolute return mandates.

Each of these stages has distinct characteristics, and will have different utility payoffs.

5.2 *Types of Investment Managers*

5.2.1 In Hodgson *et al.* (2000) the authors set out the range of manager types to consider, as follows:

- (1) *Passive investment managers.* This manager type comprises index-tracking managers. These are investment managers who do not make active investment decisions, and whose objective is to track closely the performance of a specified index. Passive management offers stable and consistent relative returns, at a low level of active risk. The lower fees charged by passive managers help to reduce the ongoing costs, and may enhance the net information ratio at the total fund level.
- (2) *Active investment managers.* Active managers apply various types of judgement to the selection of portfolios, with the objective of outperforming a benchmark. To date, the majority of assets have been invested in this way. This investment manager type can be divided into two different sub-components:
 - multi-asset (balanced) mandates; and
 - specialist mandates.

Active management offers the potential for large active returns, although the successful selection of active managers is difficult. It is possible to divide active managers into two sub-groups, active core and active satellite. Active core managers tend to be mainstream managers, who operate portfolios with wide diversifications and low active risk. Satellite managers, typically, take more risk, and, depending on the skill of the manager, can achieve higher information ratios.




Figure 6. Types of investment managers

- (3) *Absolute return managers.* Absolute return investment managers follow investment strategies that are generally less liquid, and hence long term in nature, with no explicit benchmark. This investment type contains numerous vehicles, but the main types are private equity and hedge funds. The argument for the use of the absolute return type stems from the potential for considerable performance, combined with diversification benefits.

5.2.2 We develop the application of risk budgeting to manager structure and selection in Section 5.2.3.

5.2.3 *Example 4: risk budget; policy and implementation risks*

5.2.3.1 This example looks at the risk return trade-off of a particular manager structure. While, normally, an evaluation of the manager structure might be limited to the relative return framework, the additional evaluation in A/L terms usefully yields another dimension to the manager structure decision.

5.2.3.2 The key assumptions for this assessment include manager/manager type returns, their volatilities, and the correlations (all relative to the appropriate benchmark).

Table 5. Example 4: policy and implementation risks

Relative to FRS 17	Baseline strategy	Manager structure	Combined
Expected returns			
Out performance (% p.a.)	2.0	0.2	2.2
Tracking error (% p.a.)	13.6	1.3	13.7
Information ratio (%)	14.9	14.9	16.3
Values (£m p.a.)			
Expected gain	20.2	2.0	22.2
Possible loss (VaR)	224.4	22.2	225.5

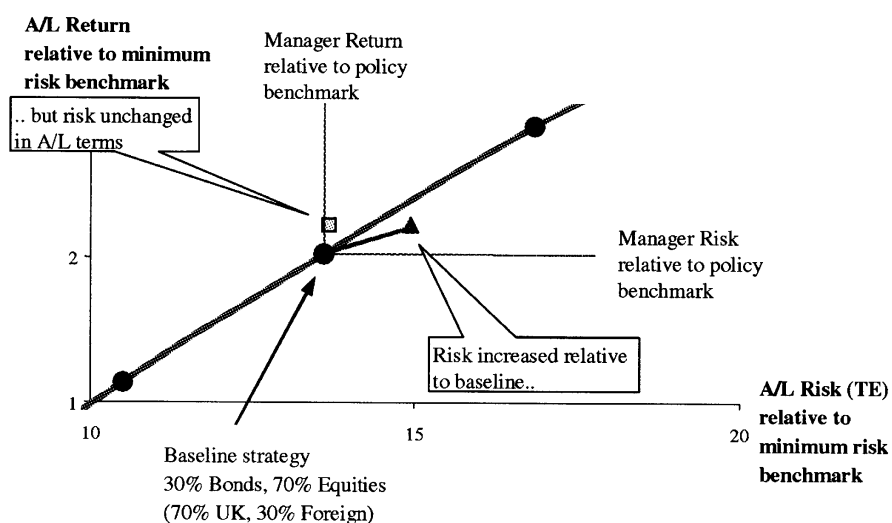


Figure 7. Example 4: risk and return trade-off

5.2.3.3 In the example, the information ratio is the same for policy and implementation risks. By combining the two activities, however, a higher overall information ratio can be produced, as the two elements are assumed to be approximately independent.

5.2.4 Example 5: risk budget; combination of strategic enhancements

The individual investment strategy decisions taken in the four preceding examples can be combined in an overall result. The low level of correlation between these sources of risk means that, while the alphas add, the risks do not.

Table 6. Example 5: combination of strategic enhancements

Relative to FRS 17	Baseline strategy	Example 2	Example 3	Example 4	Combined
Expected returns					
Out performance (% p.a.)	2.0	2.1	2.3	2.2	2.5
Tracking error (% p.a.)	13.6	13.2	13.4	13.7	13.1
Information ratio (%)	14.9	15.6	16.9	16.3	19.0
Values (£m p.a.)					
Expected gain	20.2	20.6	22.6	22.2	25.0
Possible loss (VaR)	224.4	218.6	221.5	225.5	216.4

Note: Combined strategy is 50% foreign, 5% private equity, and employs active management.

5.3 Combined Risk Budgets

5.3.1 There are two main points to note. First, the risk levels generally taken in strategic asset allocation (policy) are of the order of ten times larger than those in manager structure (implementation). Secondly, as policy and implementation risks are approximately independent, their combination produces a lower level of total risk than that implied by the sum of their respective risks. A tentative conclusion, therefore, is that any gains from manager structure after fees are worthwhile, although such gains should be evident *ex ante*.

5.3.2 These observations raise an important question in terms of where the risk budget is best spent. Should funds employ more risk in the policy (strategic asset allocation) or in the implementation (investment management structure)? The conclusion to this problem lies both in the assumptions adopted and in the governance employed. We comment on this issue in the following section.

6. PRACTICAL ASSUMPTIONS IN RISK BUDGETING

6.1 Problems: Fuzzy and Non-Fuzzy Return Estimates

6.1.1 Risk budgeting is heavily reliant on assumptions, but, in practice, all assumptions are subject to considerable estimation error. Estimates for returns from equities and bonds are subject to a lower level of estimation error, because a substantial historic data set exists. These assumptions are less 'fuzzy', as the historic data allow you to produce a relatively more reliable central estimate.

6.1.2 In most other areas estimates are subject to more estimation error, as there are insufficient data, giving a less accurate central estimate. These problems are, perhaps, most significant in setting assumptions for active manager returns.

6.1.3 We would argue that relatively non-fuzzy assumptions can be

made for U.K. equities, foreign equities, bonds, and passive manager returns. On the other hand, for alternative assets and returns of all other manager types, the assumptions will necessarily be highly fuzzy.

6.2 *Assumptions for Strategic Asset Allocation*

6.2.1 Assumptions for mainstream assets (bonds and equities) can be derived from historic analysis. Whatever historic period is chosen represents an incomplete sample. This suggests a 'Bayesian' style adjustment, that weights the prospective return as:

$$\begin{aligned} & (\text{weight1}) \times \text{capital market efficiency 'prior'} \\ & + (\text{weight2}) \times \text{historic average 'heuristic'} \end{aligned}$$

where (weight1 + weight2) sums to 1.

6.2.2 For mainstream asset classes, where good data exist, we might have more confidence in the information content of this past data, and, therefore, the weights attached to the prior can be lower than the historic average.

6.2.3 Setting assumptions for alternative asset classes is, in principal, the same. However, the weight of the 'prior' should be much greater, to reflect the fact that the experience (historic data) is much more limited. There are also technical problems, such as 'survivorship bias' and 'data mining', associated with introducing new asset classes, which, inevitably, are being considered, because their recent performance has been good.

6.2.4 We also need to be careful of placing a different emphasis on the most recent performance for different asset classes. For example, this could give a greater emphasis to recent favourable conditions for different asset classes, possibly generating an over-optimistic case for some asset classes.

6.3 *Assumptions for Investment Manager Structure and Selection*

6.3.1 Similar principles apply to setting assumptions for investment manager structure, but the problems are greater. In this area there are much greater degrees of fuzziness, and there are generally false levels of optimism, based on a misunderstanding of the influence of regression. See Urwin (1998) for a description of the 'factor of five phenomenon', the hypothesis that manager results are distributed with five times the spread of the underlying prior skill that generates them.

6.3.2 Some mandates in certain asset classes may, however, support higher overall information ratios and higher 'best in class' information ratios. Building assumptions in manager structures, therefore, involve regard for both the mandate and the manager. In addition, the level of the governance will also be a factor, with many trustee boards having structures that make successful choice of managers problematic, and which limit the level of conviction to active investment management styles in difficult times.

6.3.3 Ultimately, the subjectivity of these assumptions makes belief




Figure 8. Belief systems

systems a relevant consideration. Figure 8 illustrates how belief systems can be incorporated into the process.

6.3.4 The strength of the belief of the governing board in the differential between 'good in class' and weaker managers supports differing information ratio assumptions for different managers. Furthermore, the strength of the belief of the board in its own ability to identify and appoint outperforming managers, and to terminate the appointment at the right time, can, and should, influence the assumptions used.

6.4 Mathematics of Investment Manager Structure Risk Budgeting

6.4.1 Grinold & Kahn (2000) set out the two sources of superior information ratio as *insight* and *breadth*. Insight describes the ability to invest with superior information or judgement, while breadth refers to the opportunity to invest over a larger pool of investment opportunities. Insight is naturally a relative term, that measures the investment skill of the investment manager relative to both the efficiency of the index and the peer group of investment managers, who must, collectively, try to beat the index.

6.4.2 In the context of a pension fund's total fund, there is a third dimension to be considered in understanding superior information ratios, namely *diversification*. In this context, we define diversification as investing in a pool of investment opportunities less correlated with other pools used.

6.4.3 It is worth noting the interplay between insight, breadth and diversification. An investment manager who possesses skill, and therefore insight, in one segment of the market might not be able to generalise this skill to a wider portion of the market, or to other markets. Therefore, seeking to

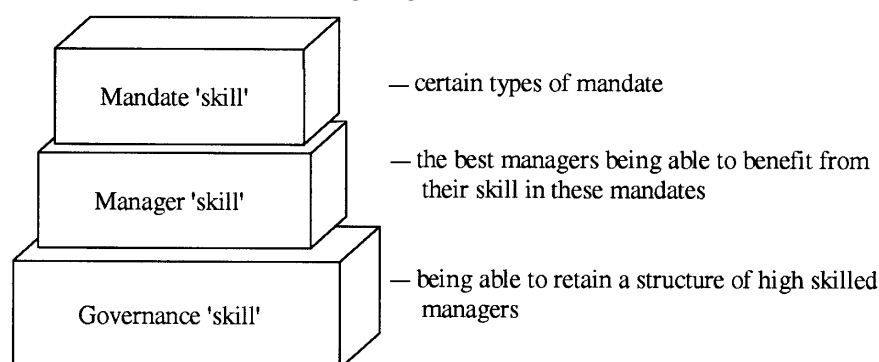


Figure 9. Building block approach to assumptions

improve the information ratio by broadening a manager's mandate might not result in improved investment efficiency, because the greater breadth could be more than outweighed by the lesser insight.

6.4.4 The alternative is to adopt a range of investment managers with strong insight (high skill) in relatively narrowly defined, but not overlapping, investment briefs. This should offer the possibility of high skill across a broad investment pool. Furthermore, as long as the aggregation of the individual investment manager mandates is constructed to maintain and, indeed, maximise the diversification element between markets and risk factors, the combination of managers with high skill and mandates with an overall broad investment span will ensure high investment efficiency at the total fund level.

6.4.5 These arguments support a building block approach to assumptions, which needs to incorporate the three key views as to the sources and sizes of superior information ratio results. An illustration of this approach is found in Figure 9.

6.4.6 On a practical point, we conjecture realistic good-in-class (gross) information ratios, for individual managers, as figures in the range 15% to 30%.

6.5 *Where should Risk be Taken?*

6.5.1 We return to the central question: "Where should funds take their risks, and in what proportions?" We tackle the question first from an empirical perspective.

6.5.2 The two fundamental levels of risk are:

- *policy risks*; strategic asset allocation and benchmark design; and
- *implementation risks*; investment manager structure and manager selection.

6.5.3 Previous studies have identified the proportions of risks taken by United States pension funds in these two areas:

- Brinson, Hood & Beebower (1986) calculated that, on average, 90% of risk *over time* is policy risk, 10% is implementation risk.
- Ibbotson & Kaplan (2000) calculated that 40% of risk *across funds* is policy risk, 60% is implementation risk. In the same paper, they found that, on average, 100% of *return* levels across funds are derived from policy risk, 0% is derived from implementation risk.

6.5.4 We can infer that:

- (1) Policy risk has the most potential performance influence for the typical pension fund. This is illustrated in Example 4, in which the value at risk from policy is a little more than 10 times the implementation risk, as mentioned in ¶5.3.1.
- (2) However, policy risk from one fund to the next is generally employed in a narrow range of asset allocations, making it, in practice, slightly less important than implementation risk when comparing risk budgets between funds.
- (3) Implementation risk produces no positive return across the industry, on average, while policy risk does have a positive return.

6.5.5 Putting these research findings into the framework created in this paper, we can deduce some interesting findings relating to the average views of the pension fund market with regard to policy and implementation risk premiums.

6.6 Risk across Funds

6.6.1 In many markets, pension funds tend to set policy in light of their peer group of similar pension funds, and, as a result, engage in an assessment of risk and return opportunities across a fairly narrow spectrum of policy options. In Table 7, using the Ibbotson & Kaplan (2000) findings, we infer the implied ratio of implementation IR to policy IR, based on actual peer group behaviour among U.S. pension funds.

Table 7. Comparison of implied implementation IR to policy IR across funds

	Assumed correlation between implementation risk and policy risk		
	0%	20%	40%
Implied implementation IR relative to policy IR of 15%*	22%	19%	18%
Ratio: implied implementation IR to policy IR	1.5	1.3	1.2

* The implementation IR required to make the 40/60 split of total risk budget between policy and implementation appears optimal.

6.6.2 We see that the majority of funds implicitly believe that they are able to add more value per unit of risk through manager structure and selection than they can through policy, when viewed cross-sectionally. The widely accepted view, however, is that, across the industry, implementation risk is essentially unrewarded. This suggests that a strong behavioural bias is at work, namely that pension fund trustees are overconfident in their belief that they can formulate investment structures that beat their peers.

6.7 *Risk over Time*

6.7.1 The story is different when viewed from the point of view of ascribing the risk over time. Here the Brinson, Hood & Beebower result, that 90% of investment risk over time is attributable to policy, suggests the results in Table 8 for the average implied assumption for the ratio of implementation IR to policy IR.

Table 8. Comparison of implied implementation IR to policy IR over time

	Assumed correlation between implementation risk and policy risk		
	0%	20%	40%
Implied implementation IR relative to policy IR of 15%*	1.7%	4.5%	7.3%
Ratio: implied implementation IR to policy IR	0.11	0.30	0.49

*The implementation IR required to make the 90/10 split of total risk budget between policy and implementation appears optimal.

6.7.2 Looking at the performance sources for a pension fund against its liabilities, the average assumption inherent within the pension fund industry seems to be that implementation adds a very small incremental amount of alpha compared to the risk involved. This seems to us to be a fair reflection of the fact that, on average, pension funds tend not to add value through their investment manager implementation. The conclusion, therefore, is that, for whatever reason, the manner in which risk budgets have been apportioned for the average fund in the industry appears reasonable.

6.8 *Conclusions*

6.8.1 What conclusions can we reach on where risks should be taken? We suggest that the governance of the fund is an essential input to any decisions, and that the principles of risk taking should be conditioned on the level of the available governance budget:

- (1) Under normal governance, low implementation risk is likely to be efficient, which corresponds to relatively simple manager structures and large allocations to passive management.




Figure 10. Governance, risk and return trade-off

- (2) Under higher levels of governance, higher implementation risk is likely to be efficient, particularly so with allocations to alternative asset classes.
- (3) Under all levels of governance, policy risk must be taken in some quantity, but this does not imply that the same type of policy risk is appropriate for all. For higher levels of governance, certain aspects of benchmark design, such as reduced home bias and higher corporate bond exposure, provide better trade-offs than taking pure equity risks. These sources of policy risk and return are not readily available to the lower governance funds, implying a shift in the focus of policy risk as the governance budget is increased.

6.8.2 The risk budgeting approach suggests that the key process of improving efficiency is by assessing marginal contribution to the information ratio (MCTIR). Policy and implementation risks that contribute most to efficiency will have the largest positive effect on IR in terms of their (marginal) incremental contribution. The methodology is similar to the approach used in the construction of index tracking funds, using BARRA or other risk measurement software. The difference, in this context, is applying the technique to all sources of both risk and return for a fund relative to a common definition of the fund's liabilities.

7. CONCLUSIONS

7.1 In this paper we have extended the concept of ‘investment efficiency’ from investment management structures to include strategic asset allocation. We conclude that risk management and measurement are vital. Risk budgeting represents a valuable way of incorporating risk and return information to produce more efficient investment decisions, and also to monitor the results of such decisions.

7.2 Risk budgeting leads to certain conclusions as to where risks should be taken, in priority order, and forms an inclusive model in which the interaction of sources of risk are explicitly considered. In its purest form, the risk budget should be allocated, based upon the marginal contribution to the information ratio for different sources of risk.

7.3 The decisions required to formulate a strategic asset allocation policy, to design an investment management structure and to appoint investment managers are inherently complex. The interplay of financial and non-financial factors often serves to further obscure the appropriate course of action. Improved governance provides the means by which these issues can be properly addressed.

7.4 We conclude, in summary, these best practice principles:

- (1) Funds should aim to increase their governance budget through better resourcing and organisational design.
- (2) Funds should consider their non-financial requirements, but aim to limit their influence.
- (3) Funds should aim to maximise the net performance per unit of risk viewed in asset/liability measures, through the creation of an appropriate risk budget.
- (4) Funds should implement their risk budget through a strategic asset allocation (policy) in which:
 - bonds provide safety;
 - equities provide return; and
 - alternative assets provide low correlation and/or additional return;

and through investment manager structures (implementation), in which:

- passive managers provide low cost safety and operational benefits;
- active core managers provide alpha under lower governance; and
- satellite managers provide greater alpha subject to greater governance.

ACKNOWLEDGEMENTS

In preparing this paper, the authors wish to acknowledge the considerable contributions of Alen Ong, Hannah Raines and Gillian Emly.

REFERENCES

- BRINSON, G.P.L., HOOD, R. & BEEBOWER, G.L. (1986). Determinants of portfolio performance. *Financial Analysts Journal*, May/June 1986.
- GRINOLD, R.C. & KAHN, R.N. (1999). *Active portfolio management: a quantitative approach to producing superior returns and selecting superior returns and controlling risk*. McGraw-Hill.
- HODGSON, T.M., BREBAN, S., FORD, C.L., STREATFIELD, M.P. & URWIN, R.C. (2000). The concept of investment efficiency and its application to investment management structures. *B.A.J.* **6**, 451-545.
- IBBOTSON, R.G. & KAPLAN, P.D. (2000). Does asset allocation policy explain 40, 90 or 100 percent of performance? *Financial Analysts Journal*, January/February 2000.
- URWIN, R.C. (1998). Avoiding disappointment in investment manager selection. Paper presented to the International Association of Consulting Actuaries, March 1998.

APPENDIX

GLOSSARY

Active return	The return on a portfolio or fund relative to a stated benchmark (also known as alpha and implementation return).
A/L return	Total return relative to liabilities.
A/L risk	Standard deviation of return relative to liabilities.
Benchmark	A performance yardstick which is investable, and against which the investment performance of an investment manager can be compared for the purposes of determining investment skill. The benchmark has a dual use in representing a strategic asset allocation policy.
Comparator	A performance yardstick which is not investable.
Governance	The organisational structure and style by which trustees (or any other governing body) carry out their investment responsibilities.
Governance budget	The level of knowledge and time resource that a trustee group, or other governing body, has in order to manage its investment management arrangements.
Information ratio (IR)	The annualised alpha divided by the annual standard deviation of alpha. It measures the significance of the alpha, and applies to both implementation returns and A/L returns.
Investment management structure	The framework which establishes how investment assets should be divided amongst different investment approaches and different investment manager types. The investment approaches can encompass different expected risk, return and style characteristics.
Net information ratio	The information ratio after allowing for costs.
Policy	Long-term asset allocation set with reference to a fund's liability profile, cash flow and funding level, and using long-term assumptions of expected future return in each asset class. It does not take into account short-term pricing anomalies between asset classes (also known as strategic asset allocation).
Policy return	The additional return relative to a baseline (e.g. the FRS 17 basis) achieved by varying the mix between asset classes.

Policy risk	The additional risk relative to a baseline (e.g. the FRS 17 basis) undertaken by varying the risks between asset classes.
Regret risk	The fear of regret following a decision that proves unsuccessful.
Risk budget	The term commonly used to describe the amount of tracking error that a fund is prepared to accept relative to a baseline (e.g. the FRS 17 basis) (also known as risk tolerance).
Theta factors	The non-financial payoffs that trustees and other governing bodies derive from certain investment positions and decisions.
Tracking error	The volatility of the alpha relative to the benchmark, calculated as the standard deviation of alpha (also known as active risk, implementation risk or sigma).