Morgan Stanley

MORGAN STANLEY RESEARCH

EUROPE

Morgan Stanley & Co. International Limited

David K Miles

David.Miles@morganstanley.com

+44 (0)20 7425 1820

Melanie Baker, CFA

Melanie.Baker@morganstanley.com

+44 (0)20 7425 8607

April 3, 2006

Economics

A Framework for Making Pension Fund Investment Decisions

Those who make decisions about asset allocation for DB pension funds face tough choices. We develop a set of tools to aid thinking about those decisions. Our aim is not to give a precise recommendation on portfolio allocation and investment strategy that can be applied mechanically to any pension scheme. Rather it is to develop a framework for thinking about the range of factors that are relevant and how they should affect investment decisions.

We ask: *what do the economics say would be the right investment decisions when there is an obligation to pay a pre-defined benefit but when it is not possible to hedge against all the relevant risks?* The way to use the model we develop is to follow two steps:

1. Figure out what acceptable risks are and how painful having assets fall short of pension costs in the future is. We provide a way of helping with this process.
2. Decide on what the economic environment is — in terms of strength of sponsor, current funding level, variability in asset returns and scale of longevity risks.

Once we have specified 1 and 2, the model gives an idea of what the right investment strategy is.

We find that under a fairly wide range of relevant conditions, a substantial degree of hedging against interest rate risk is likely to be desirable. The simplest interpretation of this result is that long-duration bonds should form a substantial part of pension portfolios — more than is currently typical. But the same end can be achieved by less direct means — through swaps and other derivatives.

This result is not driven by pension regulation or accounting rules. The question we pose is what is the right portfolio of assets based on economic factors, and in isolation from any constraints from regulations and accounting. That is significant because many view pressures on trustees and advisors to hold more bond-like assets as pushing them towards decisions that, on economic grounds, are not very sensible. Our results suggest that this is not generally the case.

But we also show which types of pension fund should do much less than the average degree of interest rate hedging.

A Framework for Making Pension Fund Investment Decisions

Introduction and Summary

What should defined benefit (DB) pension fund investments look like? Trustees and advisors face a tough set of decisions and, because of the size of existing liabilities, the decisions they make have potentially major consequences for asset pricing.

These decisions are difficult because the right structure of the assets depends on many factors. Both pension fund liabilities and assets are harder to forecast the longer ahead one looks. Deciding what the right goal should be is also difficult. Simple statements of what a trustee's duty is are rather unhelpful in deciding on portfolio allocation.[1]

In this report, we set out a framework for thinking about those investment decisions. What we do is analyse — in a simplified and abstract setting — how the many factors that affect the future balance between assets and the cost of paying pensions should influence investment decisions today. We take account of many sources of uncertainty: the pattern of future asset returns, longevity risks, the strength of the corporate sponsor. Crucial factors are also the levels of current funding and the asymmetry in how upside and downside risk to fund performance are regarded. We hope that the tool we have developed might help in thinking about investment decisions for DB pensions.

We develop a calibrated model which — given assumptions about key inputs — generates explicit answers about the optimal structure of assets. The model is only a rough guide to what any particular pension fund should do. We see it more as a tool to help the thought process behind investment decisions and not something that should be used mechanically to produce answers about investment strategy. Although it is a stylised and somewhat abstract model, we feel it could help guide the thinking of those who need to make investment decisions. Those scheme-specific decisions will of course be based on detailed information about plan structure, member demographics and the current and potential scale of contributions from sponsors and plan members. What the analysis we describe here can do is help shape some of the thinking behind those detailed decisions.[2]

[1] For example, to say that they should act prudently and responsibly to help ensure that the pensions are paid is too vague. What if there is a strategy that means that there is a 90% chance pensions are paid in full, but a 10% chance there is only enough to pay 80% of liabilities. Is that better than an alternative strategy that gives an 80% chance of being paid in full and a 20% chance of pensioners getting 90% of what they expect?

[2] This is in the same spirit as the aims of the analysis undertaken by Haberman and colleagues (2003) and by Orszag and colleagues (2005). For full references see the appendix.

Key features of the results

Under a wide range of scenarios, we find that a substantial hedging of shifts in yield is likely to be optimal. In the simplest interpretation of this result, it means that long-duration bonds are likely to form a substantial part of pension portfolios. But the same end can be achieved by less direct means — through swaps and other derivatives.[3]

This finding does not reflect pension regulations or accounting rules. In fact, the question we pose in our work is: what is the right portfolio of assets based on economic factors, and in isolation from any constraints from regulations and accounting? We find that the answer to this question — under many conditions — is that more hedging against interest rate risk should probably be done than is currently typical.

This is significant because many view pressures on trustees and advisors to hold more bond-like assets as pushing them towards decisions that on economic grounds are not very sensible. In general, we find that this is not the case.

But the appropriate asset allocation will differ across funds. The model suggests that there are a range of factors that should be important for investment decisions. Some of the key results are:

- The volatility of returns on equities (our shorthand for alternatives to bonds) matters a great deal. One would probably need to believe that such volatility was a good deal lower than has been typical in the past to warrant equity weights as high as 60%.
- More longevity risk — which we assume cannot be hedged, and certainly not with bonds — does NOT mean that a more conservative investment strategy is optimal.
- Whether bond yields are likely to rise (or fall) may be more important than whether they are low (or high) relative to returns on other assets.
- The right investment strategy may not be very sensitive to the current funding level when the appropriate horizon for judging performance is long.
- Risk tolerance is central to investment decisions. But quantifying risk appetite is tricky because the upside and downside risks for trustees are so different. We believe that our way of calibrating attitudes is useful here.

[3] We will in this report often rather loosely refer to allocations towards bonds, or to bond-like assets. We use this as a shorthand way of referring to a strategy of having the assets structured so that their value moves in a way more closely related to shifts in longer-term bond yields.

What the model does…

One of the reasons why making investment decisions for pension funds is very hard is that there are multiple sources of uncertainty (including longevity, the future health of the plan sponsor and asset price volatility). These sources of uncertainty become larger the longer the investment horizon is. We believe that it is very useful to try to specify as clearly as possible what the trade-offs between different types of risk are. Once the trade-offs are clear, we can better assess one of the most important and thorny issues facing pension funds today — how far should they go down the road of matching debt-like liabilities with assets that share those characteristics. Potentially, that strategy comes at some cost — it cannot reduce all risk and may mean lower exposure to other asset classes which one might expect to yield more, on average, but be both more volatile and less correlated with the market value of liabilities.

The central issue we explore is what is the right amount of hedging against interest rate (i.e., yield) risk. (We rather loosely refer to this throughout the report in terms of 'bond' allocation.[4]) We are focused less on the best way to implement hedging against yield risk than the more fundamental question of what should drive the right amount of interest rate hedging. Central to this question is the fact that there are risks that cannot be hedged easily — and certainly not with debt-like instruments.

We think that having a clear and explicit model to address this issue is very useful — even if it is necessarily abstract and has to simplify from many real world factors. It is *not* intuitively obvious what the impact on portfolio choice should be of changes in current funding levels, the scale of longevity risk, the volatility of bond yields or in the chance of sponsors being unable to bail out pension funds in deficit. A tool that gives you some quantified answers to that (albeit one based on a stylised version of reality) has potentially great value.

We see two obvious uses of the model we have developed and the results that come from it: First, as a guide to people who have to make investment decisions about pension fund assets. Second, as a means of assessing the gap between the current structure of typical pension fund assets and what might be optimal. The scale of that gap is likely to be a key driver of investment flows and potentially have significant asset price implications. At the moment, a typical DB pension fund in the UK and the US has an equity weighting of just above 60%, a bond weighting of slightly below 30% and an allocation to property, cash and alternative assets that is, on average, relatively small (around 10% in total). Roughly how far from the right structure of assets, *based on the fundamental goals of pension funds* (and not on the secondary issue of how they are accounted for, nor even how they are regulated), is that? Given the large stock of assets in UK pension funds — of the order of £800 billion, or around 70% of GDP — we believe that the answer to this question is of great significance for the demand for various asset classes, and therefore potentially for asset pricing.

The issues we address in this report have a much wider significance than just in the UK. The problems facing those that control pension fund assets are similar in many countries: US and European DB pension schemes are also of great economic significance. The question of the right way to structure the portfolio assets in a world with many sources of uncertainty, where hedging all risk is impossible and when many schemes start off in deficit, is as relevant in the US and in many other European countries as it is in the UK.

Plan of the report

In the next section we outline the issues we address in more detail and describe the structure of the model. We keep technical details to a minimum. We then describe the results on optimal pension fund asset allocation in detail. In the final section of the main report, we consider the market implications of the results — both at the micro (fund) level and in aggregate.

An appendix provides full detail of the model. More information is available on request.

The Pensions Landscape and an Overview of the Model

In the UK and the US, more and more DB schemes are closed to new members. In the UK, only a minority of private-sector DB schemes are open to new members. For all UK DB schemes, roughly one-third of members are active (still contributing and accruing liabilities) but two-thirds are either deferred members or members already receiving pensions. Given that current position — and the likelihood of more DB schemes closing to new and existing members — the investment decisions made about the portfolio allocation of existing assets will become ever more important relative to the significance of new contributions coming into funds.

[4] That hedging can be done in lots of ways, e.g., through derivatives (such as swaps), and not just through direct holding of long-duration bonds. But ultimately it is likely to create demand for long-duration assets.

We focus on asset allocation and consider how a DB fund should structure its assets to meet future pension payments, which are liabilities already accrued. In doing this, we allow for the (uncertain) possibility of top-ups to the fund from the sponsor if assets fall short of liabilities as a scheme matures.

The range of factors we incorporate in our model is lengthy. Some are specific to individual pension schemes and some are common. The key factors include:

- Likely future level and volatility of bond yields;
- Scale of longevity risks;
- Expected return and volatility on assets other than bonds (which we call 'equities', but that is just a shorthand for a broader range of alternatives);
- Strength of sponsoring company — measured as the probability that the sponsor is willing and able to make good any funding shortfall in the future;
- The existing funding level of schemes;
- Attitudes to risk of the scheme managers — who we will call the trustees. A key aspect of this attitude to risk is the degree of asymmetry between the pain of losses and the gain from good performance of the assets in the fund.

The nature of the pension fund decision

The way we analyse the investment decision is based on a simple but fundamental assumption: those with the responsibility for asset allocation consider what the implications are for the possible future deficit/surplus of the fund.[5] We focus on the decisions to be taken now about the appropriate allocation of assets between two broad classes of assets. We refer, as a shorthand, to these asset classes as 'bonds' and 'equities'.[6]

We want to focus on allocating existing assets, so we abstract from future pension contributions and from the accrual of new pension liabilities. But we do allow for the possibility of a sponsoring company making good any shortfall of asset values over the future value of existing pension liabilities (i.e., making good future pension deficits). These assumptions are obviously major simplifications. But we believe that the nature of the results we derive is nonetheless relevant for relatively mature pension funds (which make up a large and rising proportion of DB funds)

To be more specific, the pension problem we analyse is one where a decision is made today about allocating a portfolio of assets where the results of that decision will be known several years down the line. At that time, the asset values that result will be judged against the value of the pension liabilities at that point. Both asset values and pension liabilities in the model are uncertain. Some of the factors that influence them are common and affect asset prices (e.g., bond yields) while others are assumed to be quite distinct (e.g., longevity has a major effect upon the value of pension liabilities but we assume that it does not impact the value of any assets).

We think of the relevant investment horizon as the period up to the point at which the value of pension liabilities, on average, gets realised. For the pension liability of an individual worker, one might think of that as the period until they retire.[7] We will value the fund's liabilities at this point using the cost of an annuity, given bond yields and life expectancy at that date. Of course, in practice pension funds have a wide range of types of liability (with many pensions already in payment as well as accrued liabilities to current active and deferred members of various ages). In our calibrated model, we set a single horizon (which we think of as the weighted average time from now until pensions are first taken). We set a plausible value for that time by reference to both the age structure of actual members of UK pension schemes and also by reference to the duration of the liabilities which it generates (which also depends on life expectancy and bond yields).

It is a major simplification to assume that investment decisions made today are not changed in the future — so the results we derive are at best an approximation to a truly dynamic optimal investment strategy. They are a guide to the direction in which asset allocation might move, rather than a prescription for a change to be made at the first opportunity.

[5] This does not mean that those responsible for investment decisions ignore the interests of the sponsoring company and its shareholders. In our framework, the prospect that asset outperformance can generate a fund surplus can be valued, though most of that upside benefit is likely to accrue to the sponsoring company.

[6] We stress once again that when we talk about the bond weight in the portfolio, this is just a shorthand way of describing the extent to which the portfolio value moves in line with yields, or bond prices. (In practice that can be achieved with derivatives and swaps and not just through holding bonds.) And when we talk about 'equities', this is a shorthand for assets whose returns are not highly correlated with bonds and may offer (on average) higher returns, but have more variability. In other words, 'equities' is shorthand for a composite alternative asset class to bonds — one with a higher average yield, but which is less correlated with the value of liabilities.

[7] At which point the liability could be extinguished and an annuity bought. Even if it is not extinguished, its value at that point becomes much more certain.

The Pension Model and Sources of Uncertainty

How the model works

We assume that pension fund assets can be invested in assets with values linked to bond yields ('bonds') and alternative assets ('equities').

The liabilities are commitments to pay pensions in the future. We define a horizon (T) which we think of as the average time to when those pensions are first paid. At horizon T a commitment to pay a pension which gives a constant real flow of income has a value that we model as proportional to a level annuity. The value of that annuity at time T, and hence the value of the liabilities, depends on two variables that are highly uncertain — real yields on bonds and life expectancy (both at time T). We make the simplifying assumption that the slope of the yield curve is constant and all uncertainty is about the level of real yields.[8]

In the model we assume that liabilities are commitments to pay real (inflation-adjusted) pensions. This is why the relevant returns are always real returns — hence our focus on the real yield on bonds, and uncertainty about it, and the real yield on other assets.

Real bond yields at time T are, when viewed from today, uncertain: they are assumed to be equal to yields today plus a 'shift factor' that reflects anticipated changes in yields and a random element. The random element can be large relative to yields today. We describe how we calibrate its magnitude below.

The return on 'equities' is also uncertain. It depends on an expected annual rate of return plus a random element (where log returns and the random element are assumed to be normally distributed).[9] We also consider the effect of correlation between bond yields and equity returns, though in practice we find that this is not a very important factor.

The pension fund has to choose an allocation to 'bonds' and 'equities' given the current value of assets. A key question is by what criterion are those decisions judged.

By what criterion do we judge outcomes?

We assume that what matters to a pension fund (or to its trustees and advisors) is the balance between assets and liabilities at horizon T. That balance depends on a wide range of factors in our model: it reflects portfolio decisions made today, how life expectancy evolves; how the value of bond yields and equity values has moved between now (time t) and T; and the possibility that a sponsoring company makes good any shortfall in asset values at time T.

Crucial to the optimal choices at time t is how those responsible for managing the fund (who we will from now on call the 'trustees') assess the costs of having a deficit at date T, and what the benefits (if any) of assets exceeding liabilities at T are. Those costs, and potential benefits, are likely to be highly asymmetric. In other words, the cost of assets falling 10% short of liabilities at date T is likely to be greater than the benefit of assets exceeding liabilities by 10%. Indeed, if pensioners cannot gain any of the upside benefits of asset outperformance, *and if those managing the pension fund identify completely with the interests of the pensioners*, there may be no value in having a surplus at all. Of course, the trustees may also give weight to the interests of the sponsoring company and so attach significant value to the chance that assets exceed pension liabilities even if the entire excess may, effectively, accrue back to the sponsoring company and its shareholders.

It is very likely that the costs of a deficit at T rise more than proportionally with the scale of the shortfall. In other words, the trustees are likely to be risk averse, so that a shortfall of 20%, for example, is *more than twice* as bad as a shortfall of 10%.

Unless there is complete certainty that the sponsoring company will make good any shortfall at time T, and that any outperformance of assets over liabilities is of no value, then these risk preferences are central to the asset allocation decision.[10]

---

[8] The duration of the liability at time T depends upon both real yields and life expectancy at time T. If real yields are low, that duration is approximately half life expectancy. So one natural asset to hold today is a bond with a duration of approximately T + half the expected post-retirement life expectancy at time T. We assume that there is a zero-coupon bond with a duration equal to T plus half the anticipated post-retirement life expectancy. (The appendix spells out exactly how we value pension liabilities and assets.)

When we calibrate the model, we set T equal to 10 (years); life expectancy at age 65 now is a little under 20 years. In 10 years' time a reasonable central expectation is that it will be around 20 years. So the expected value of post-retirement life expectancy 10 years ahead is 20. This means that T plus half anticipated life expectancy is also 20 years. In fact 20 is about in line with the typical estimate of the average duration of the liabilities of UK DB pension schemes.

[9] The random element of equity returns over T years is a (normally distributed) shock with standard deviation of $\sqrt{T} \times \sigma$ (where σ is the *annual* standard deviation of equity returns).

[10] If the sponsoring company were 100% certain to make good any shortfall of assets over liabilities in the future, and also could appropriate any outperformance, then pension scheme members — and trustees acting on their behalf — should be completely indifferent to asset allocation. The shareholders in the sponsoring company might care a lot about the asset allocation. But their interests are really nothing to do with the details of the pension scheme. In fact, if shareholders can offset any asset allocation decisions in the DB fund by re-arranging their own investment portfolios, then the asset allocation of the pension fund is totally irrelevant (ignoring tax effects).

It is highly unlikely that for any DB pension fund these conditions hold, so that the risk preferences of those making decisions (who we will continue to call ‘the trustees’) about the way fund assets are invested do matter.

We describe the risk preferences of the trustees with a utility (or payoff) function that allows for both asymmetry between the effect of having a deficit or surplus, and also for more rapidly rising pain from larger deficits.

The extent to which the pain of deficits rises with the magnitude of deficits is determined by a parameter, α. The degree to which having a future surplus creates any ‘benefit’ is reflected in the value of β. If, as seems highly likely, the value of a surplus is significantly less than the pain of an identically sized deficit, then β > α.

The relative values of α and β are crucial in determining the portfolio allocation. This is just a rather jargon-heavy way of saying something pretty obvious and natural — what the right investment decisions are today depends on how painful shortfalls in asset values from the real cost of paying pensions are when pensions have to be paid. It also depends on how valuable might any outperformance be so that assets are in excess of the cost of paying promised pensions. There is no sense in which there are ‘right’ or ‘wrong’ values. Trustees face a difficult job in assessing which risks are acceptable and which are unacceptable for a pension fund. *But they cannot avoid that choice.* Given the limited ability to hedge some of the major risks that pension funds face, the option of ‘no risk’ is not available.

Without specifying the values of α and β that reflect the willingness to take risk, there is no meaning to the concept of optimal decisions. The problem is not well specified.

The way to assess what plausible values of α and β might be (and for trustees to gauge what their own values are) is to see what they imply about the acceptability of specific investment risks. We describe a way of assessing what these values mean, and what the appropriate values for pension fund trustees and their advisors are, in the following box.

Attitudes to Risk

Exhibits 1 and 2 give an indication of what various values of α and β mean in terms of the acceptability of specific risks. Exhibit 1 focuses on upside and downside risks for a scheme which is, at best, only just fully funded — this is all about the value of α. Exhibit 2 allows us to assess attitudes to risk that reflect asymmetry between upside and downsides around an initial position of being fully funded — this is all about β.

Exhibit 1 shows, for various values of α (that reflect ever-greater degrees of risk aversion), what the minimum acceptable probability of success needs to be for an investment with symmetric payoffs. Suppose we have a deficit of 10% and a degree of risk aversion of 5 (α=5). The table then shows that an investment which could, if successful, take the fund back to 100% funding, but if unsuccessful would take the funding level down to 80%, needs to have a probability of success of at least 64%. In other words, if α=5, the chance of gaining 10% (wiping out the deficit) needs to be almost twice as great as the chance of losing 10% (taking the funding level down to 80% from 90%) for that to be an acceptable risk.

For any given positive level of α, the larger the risk, the greater needs to be the chance of success for an investment to be acceptable. Based on Exhibit 1, our judgment is that the attitude to risk of most trustees is such that values of α below 2 are likely to be rare. We suspect that few trustees have a degree of risk aversion over their pension as low as even α=3. At that level, a risk that could increase a deficit from 90% to 80% — or if it goes well could make the scheme fully funded — is just acceptable if it pays off slightly under 60% of the time.

A value of α=4 is probably a better reflection of most trustees’ risk aversion. But many trustees would only feel comfortable with risks that imply a much higher value of α, so a value as high as 8 is not implausible. A value of α equal to 7 implies that a really big investment risk — one that could wipe out a 25% deficit if it pays off but would reduce the funding level from 75% to only 50% if it goes wrong — needs to have a success rate of very close to 93% to be acceptable (so that its chance of success is more than 10 times the chance of failure).

The way to read Exhibit 2 is similar to Exhibit 1. But now we focus on a different type of investment risk. We consider a risk that, if successful, takes a fund that is 100% funded to a position of being over-funded. If it is not successful, it creates a deficit of the same size. Again we focus on symmetric payoffs and consider the minimum acceptable chances of success. We show what different values of β mean for a given level of α. We take α=4.

Exhibit 1

Chance of Success for the Investment to Be Acceptable

Initial deficit and potential gain or loss (%)	Risk aversion parameter (α) 0	1	2	> increasing pain felt from deficits 3	4	5	6	7	8
5	0.50	0.51	0.53	0.54	0.55	0.57	0.58	0.59	0.60
10	0.50	0.53	0.56	0.58	0.61	0.64	0.66	0.69	0.71
15	0.50	0.54	0.59	0.63	0.67	0.71	0.75	0.78	0.81
20	0.50	0.56	0.63	0.68	0.74	0.79	0.83	0.86	0.89
25	0.50	0.58	0.67	0.74	0.80	0.86	0.90	0.93	0.95
30	0.50	0.61	0.71	0.80	0.87	0.92	0.95	0.97	0.98

Source: Morgan Stanley Research

Exhibit 2

Chance of Success for the Investment to Be Acceptable (When α=4)

Potential gain or loss (%) from a fully funded position	'Asymmetry parameter' (β) 4	6	8	> lower value placed on surpluses 10	20	40	500
5	0.55	0.56	0.57	0.58	0.64	0.72	0.97
10	0.60	0.62	0.64	0.66	0.74	0.83	0.98
15	0.65	0.68	0.70	0.72	0.81	0.89	0.99
20	0.69	0.73	0.76	0.78	0.86	0.93	0.99
25	0.74	0.77	0.80	0.83	0.90	0.95	1.00
30	0.78	0.81	0.84	0.86	0.92	0.96	1.00

Source: Morgan Stanley Research

If β=10 (when α =4) an investment that could create a surplus of 10% or a deficit of 10% would need at least a 66% chance of success to be acceptable (i.e., twice as great a chance of success as failure). At β=6, that acceptable success rate falls to 62%, while at β=20 it rises to close to 75%.

It may well be that the benefits to pensioners of a fund being more than 100% funded are perceived to be so small — and the weight attached to the sponsor getting the upside benefit so small — that no gamble that, if successful, makes a fund more than fully funded (but which if it went badly could create a deficit) is acceptable. This corresponds to very high values of β. If β is 500 or greater, then essentially no such gamble is acceptable. Such an attitude to risk would be possible if trustees identify completely with the interests of scheme members and took the view that pensioners stand virtually no chance of getting higher pensions if a fund was in surplus. For some trustees, that may be a reasonable position to take.

Once again, we stress that there are no right or wrong values of α of β. Trustees acting on behalf of pensioners need to ask themselves what the acceptable risks are to those whose future incomes are at stake. This is what α reflects. They also need to assess the chance that if schemes are over-funded, more benefits flow to pensioners, and what value to attach to the gains that would flow back to the sponsoring company (which affects the extent to which β exceeds α).

One way to judge how a trustee might assess a suitable α is by reference to typical attitudes to risk of households. There is a good deal of evidence on that. Not surprisingly, it suggests that there is great variability in risk aversion among households. But some very large-scale surveys suggest that a value of α of around 4 might reflect average risk aversion among households. In Exhibit 2 we show optimal portfolio allocation decisions when α=4; but we also show the impact of much higher and lower values.

Judging reasonable values of β is harder. There is likely to be much more diversity among trustees on this. We focus mainly on two cases below (β=∞ and β=10), but also show the impact of other values.[11]

Exhibit 3 shows the relationship between the degree of surplus or deficit at time T and the payoff to the pension fund trustees for these two combinations of α and β.

Exhibit 3

Utility ('Payoff' Function) of the Trustee

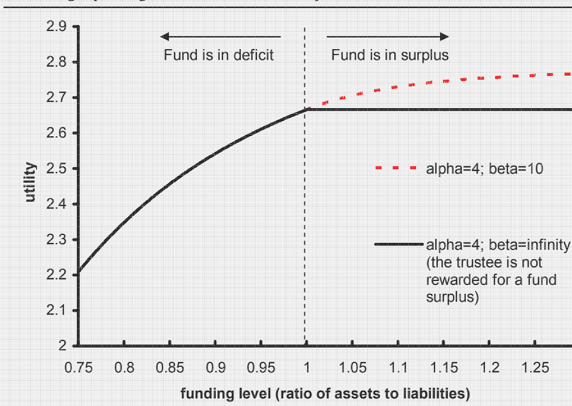

Source: Morgan Stanley Research

[11] When β = 10 we are attaching significant value to potentially having a fund surplus at the relevant horizon. This could reflect one or both of the following: that scheme members are likely to share in some of the benefits if assets outperform liabilities; or that trustees attach weight to the interests of the scheme sponsor, who gains when assets outperform.

The strength of the sponsor

A final key element of the asset allocation decision reflects the chance that a company is willing and able to make good any shortfall at time T. Clearly there is uncertainty about whether a sponsor will fill any funding gap when the relevant horizon is several years ahead. Our model assumes that there is a probability that a company is able and willing to 'make good' any deficit at time T. We assess the impact of variations in this probability on optimal asset allocation and on likely scheme solvency.

Making the Right Investment Decision

Exhibit 4 shows the information used in our framework by the trustees to make a decision on the pension fund portfolio allocation. We assume that trustees, with the help of their advisors, do that by assessing the probabilities that assets (relative to liabilities) take on various values at the T year ahead horizon, bearing in mind the probability that if the pension fund is in deficit at T the company makes good the difference. They choose the portfolio that maximises their expected utility (or 'payoff').

To solve the problem, we use Monte Carlo simulation methods. We create several thousand possible future histories using the assumed sources of uncertainty in the problem (namely, equity returns, bond yields, life expectancy and whether the company makes good any shortfall at T or not). We choose the portfolio weights that give the highest expected value of utility (or payoff) based on the outcomes for these simulated future histories. That choice depends upon attitudes to risk (i.e., upon α and β). Essentially, what we are doing is this: we assume that the trustees and their advisors take account of all the risks they face and figure out the probabilities of different outcomes for the overall balance between assets and liabilities at time T. They then take the decision that best balances risk and reward, given those probabilities, and given their appetite for risks and the asymmetry between upsides and downsides (as reflected in α and β).

In making that decision, trustees and their advisors face a range of trade-offs. If a scheme is in deficit today, they might limit the chance of a worsening deficit by having a high bond weight, but at the cost of very little chance (in the absence of a sponsor top-up) of being able to pay pensions in full. If a scheme is close to being fully funded, then taking a high bond weighting reduces the probability of not paying pensions in full but cannot remove it (because of longevity risk) and might give very little chance of assets exceeding liabilities.

These sorts of trade-off are what the model handles. The way to use it is in two steps:

1. Figure out what acceptable risks are and how painful having assets fall short of pensions costs by varying amounts in the future is. (This is about choosing α and β — see the box on page 6-7.)

2. Decide on what the economic environment is — in terms of strength of sponsor, current funding level, variability in asset returns and the scale of longevity risks.

Once we have specified 1 and 2, the model then gives an idea of what the right investment strategy is.

Exhibit 4
Parameters of the Model

Parameter	Description	Base case assumption (alternatives)
μ	Average annual rate of real return on equities	6% (4%, 5%, 7%, 8%)
ε	Random component of equity rate of return	Normally distributed, mean zero, standard deviation 17.5% (10%, 12.5%, 20%, 22.5%)
y	Real yield on bonds today	1.5% (0.75%, 1%, 2%, 3%)
y_T	Expected real yield today on bonds T years ahead	1.5% (1.0%, 2%, 2.5%, 3.0%)
σ	Standard deviation of the yield T years ahead (bp)	75bp (50bp, 100bp, 125bp, 150bp, 200bp)
T	Average horizon until pension paid (years)	10
M	Life expectancy of newly retired pensioner at T periods ahead	Equal probability of 20, 20 plus 20% and 20 minus 20% (plus or minus 5%,10%,30%)
α	Risk aversion for loss parameter	4 (1, 2, 6, 8, 10)
β	Asymmetry parameter for gains/losses around full funding	Infinity or 10 (8, 20)
f	Initial funding level (assets as a proportion of liabilities)	90% (75%, 80%, 95%, 100%, 105%, 110%)
p	Probability of a full bailout from the scheme sponsor	75% (0%, 25%, 50%, 60%, 80%, 90%)

Source: Morgan Stanley Research

The box overleaf shows the various types of economic environment we consider.

Calibration of the economic environment

The assumptions used in the model are summarised in Exhibit 4. They are as follows:

Equity returns/volatility	In the central case we set the mean real rate of return on equities at 6% per year with a standard deviation of 17.5%. These figures are close to those calculated by Dimson, Marsh and Staunton (*Triumph of the Optimists*, 2002) for world real equity returns 1900-2000 (5.8% mean and 17.0% annual standard deviation). We also consider the effect of higher and lower average rates of return on equity (of 4%, 5%, 7% and 8%) and of higher and lower annual volatility (of 10%, 12.5%, 15%, 20% and 25%).
Bond returns/volatility	In the central case, we set the real yield today at 1.5% and assume the real yield in T (10) years' time is normally distributed around the same mean level with a standard deviation of 75bp. That implies roughly a two-thirds probability that the real yield would lie in the range 2.25% to 0.75%, and a 95% probability that the real yield lies in the range 3% to 0%. The central case real yield of 1.5% accords with Bank of England data, showing that the average real yield on UK government debt across maturities since the beginning of 2005 has been 1.5%. Our own calculations of real government bond yields suggest that the standard deviation of the 10-year difference in bond yields since 1700 has been closer to 1.75% than the 0.75% we assume as the central case in the model (see *Where Should Long-Term Interest Rates Be Today? A 300-Year View*, D Miles et al, March 9, 2005). However, in the period since inflation targeting (1993-), this standard deviation has been closer to 0.75%. We also consider lower and higher mean levels of real yields today, and lower and much higher bond yield volatility. As well as this, we assess the impact of an expectation that, on average, the real yield in the future will be higher (or lower) than it is today.
Correlation	In the base case, we assume that there is no correlation between bond yields and equity returns. But we also allow for the possibility of negative correlation between the level of yields and the value of equities.
Average funding level	We take a 90% initial funding level for pension schemes in the base case. This is a little higher than the current typical funding level of UK pension schemes on an FRS17 basis (which is about 85% according to the 2005 National Association of Pension Funds survey). We also consider much lower, and somewhat higher, levels of initial funding.
Liability duration	We assume that the average time that scheme members are away from retirement is roughly 10 years. NAPF data show that for a large sample of UK DB schemes, roughly one-third of scheme members are active, roughly one-third are deferred members and roughly one-third are already receiving pensions. If the average age is around 40 for active members, 45 for deferred members and near 70 for those in receipt of pensions, the average time to a typical 62-year retirement is close to 10 years. If the expected life of retiring pensioners then is just over 20 years, then with low real yields, the duration of the pension at retirement is also close to 10. This generates an overall expected duration today of pension liabilities of 20 (the sum of the average time until the pension is first paid and the anticipated duration of the pension liability at retirement). We take the average expected duration of pension liabilities to be approximately 20 years (and therefore the natural duration of the bond portfolio held today should be 20).
Mortality of pension scheme members	We assume that the expected level of life expectancy, in 10 years' time, at retirement is 20 years. This is almost one year longer than current estimates of life expectancy for today's 65-year-old males. We assume that the risks around this central forecast are large and that the level of life expectancy could be 20% either side of the central forecast (so that there are equal probabilities that life expectancy is equal to, greater than or smaller than the central forecast of 20 years). We also consider the impact of much greater, and somewhat smaller, uncertainty over life expectancy.
Probability of company being around to bail out the fund	In our base case scenario, we assume that the probability of a company being able and willing to fund any pension deficit 10 years ahead is 75%. We also consider much lower and (somewhat) higher probabilities.
Risk aversion of trustees	Our central assumption is that trustees have a degree of risk aversion about downside outcomes (less than full funding) to make α equal 4. This is in line with survey evidence of the degree of risk aversion of households in developed economies. (For a discussion of that evidence, see *What Should Equities and Bonds Be Worth in a Risky World?*, D Miles et al, September 2005.) The acceptable risk-reward trade-off for the trustee is effectively assumed to change sharply in the case where the pension fund moves into surplus — gains above full funding in asset values are very much less valuable than gains that close a funding gap. For one set of results, we assume that asset gains above full funding create no benefits (β is infinite). We also consider a value of β of 10, which implies that there is significant value to gains in asset value beyond the point of full funding. We also explore the impact of higher and lower values of α and β.

Results

We begin by summarising results for the investment decisions where we consider a pension fund that is 90% funded and where the sponsor is assumed to have a fairly high chance of making good any shortfall that might exist at a 10-year horizon. We concentrate on two aspects of the investment decision: the degree of matching to yield risk, which we reflect in the 'weight in bonds'; and one measure of risk, which is the proportion of times the fund is in deficit (after allowing for the chance of a top-up from the sponsoring company) at time T. We present the results in a series of charts which vary some of the key factors in turn. In each chart we show the 'bond weight' and the probability of being fully funded or better at the relevant horizon.

In all cases, the right investment choice is sensitive to the degree of risk tolerance and how much, if any, value is attached to asset outperformance once we go beyond full funding. We look at two cases: one where there is some value to outperformance (β=10) and one where there is not (β is infinite).

In the central case when some value is attached to asset outperformance that creates a surplus (β=10), and when we start with a 90% funding level, the optimal bond weight is almost exactly 60%. The chances of ultimate under-funding based on that portfolio decision, and given an assumed 75% chance of a full top-up from the sponsoring company, is about 11%.

If no weight is attached to extra funding beyond the 100% level at the relevant horizon (β is infinite), the optimal portfolio allocation is significantly more conservative. The optimal weight in bonds now is 80% so that the 'equity' (i.e., non-bond) weight is half as large at 20%. But the probability of being at least fully funded is actually slightly lower. When we start from a position of under-funding it does not follow that a higher weight in bonds increases the chance of solvency.

Exhibit 5
Base Case Results

	α=∞	α=10
Bond allocation	0.79	0.60
Equity allocation	0.21	0.40
Proportion of times the fund is not in deficit at T	0.87	0.89

Source: Morgan Stanley Research

Based on our stylised model we find that the five factors which are really important for optimal asset allocation (in rough order of importance) are:

1. Volatility and average size of equity returns (see Exhibits 6-9). Holding the average real equity return steady at 6%, if volatility is lower than the central case assumption (i.e., lower than an annual standard deviation of 17.5%), the equity weight rises sharply. For example, if volatility is only 12.5%, the optimal weight in equities roughly doubles from around 20% to 40% when there is no value to having a future surplus (β infinite) and from 40% to close to 75% when β=10. Higher equity volatility (at 22.5%) has the opposite effect. The equity weight falls sharply from 40% to 24% when β=10. In contrast to the power of these effects, changes in the expected return on equities for a constant volatility of returns seem to be a less significant factor — especially when upside success is not significantly valued.

Exhibit 6
Effect of Equity Returns — No Benefit in Having a Future Fund Surplus (β=∞)

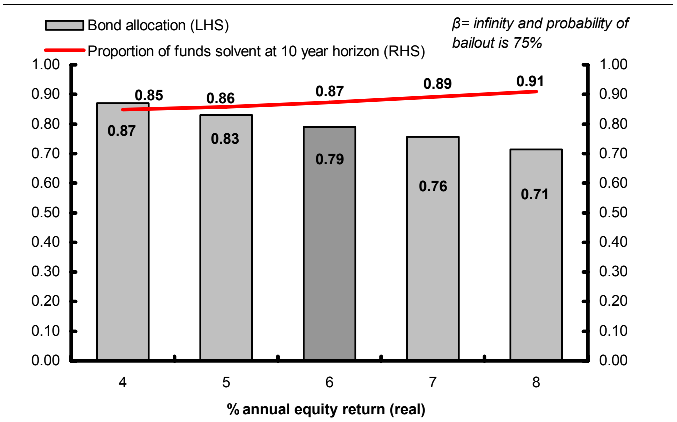

Source: Morgan Stanley Research

Exhibit 7
Effect of Equity Returns — Benefit Attached to Future Fund Surplus (β=10)

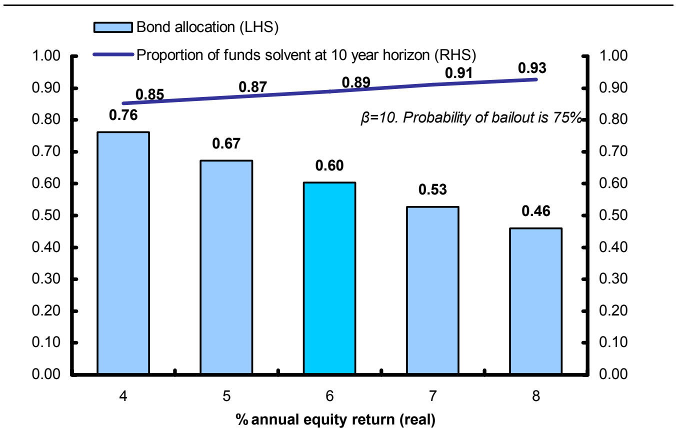

Source: Morgan Stanley Research

The key take-away here is that the volatility of returns on equities (our shorthand for alternatives to bonds) matters a great deal. One would probably need to believe such volatility was a good deal lower than has been typical in the past to warrant 'equity' weights as high as 60%.

Exhibit 8

Effect of Equity Volatility — No Benefit in Having a Future Fund Surplus (β=∞)

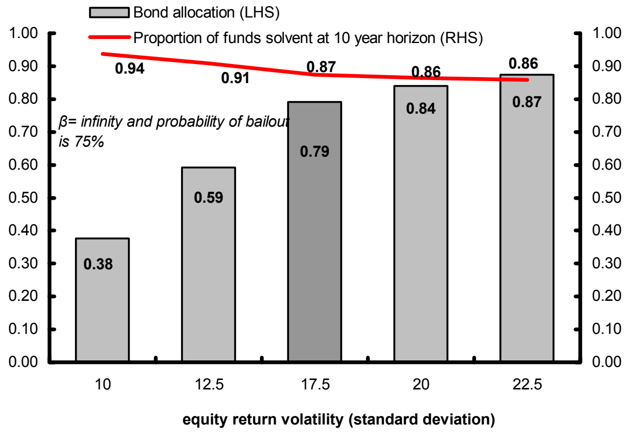

Source: Morgan Stanley Research

Exhibit 9

Effect of Equity Volatility — Benefit Attached to Future Fund Surplus (β=10)

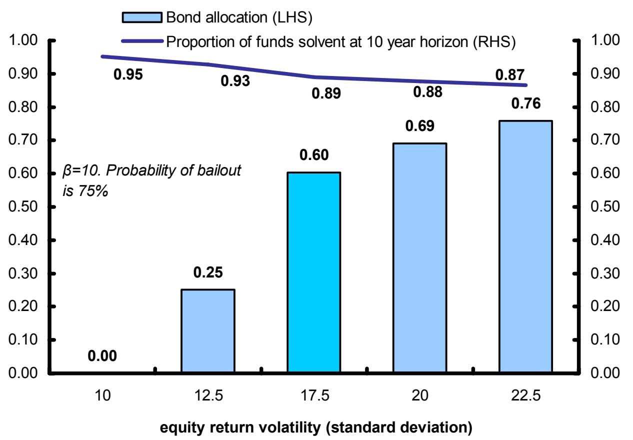

Source: Morgan Stanley Research

2. Degrees of risk aversion (see Exhibits 10-11). In the base case we assume that aversion to downside risk is such that α=4. What this means is that an investment that creates a 10% outperformance if it goes well, but creates a loss of 10% if it goes badly, needs at least a 60% chance of succeeding to be acceptable. If risk aversion is significantly less (α=2), the bond weight falls from 80% to 75% when performance beyond full funding is not valued, but falls much more, from 60% to 44%, when it is (and β=10).

But we suspect that risk aversion among decision makers is actually more likely to be greater than implied by α=4 than less. And Exhibits 10 and 11 show that if it is, then 'bond weights' should be higher. The impact can be big. Let's try to be more specific.

Suppose there is an investment risk which, if it goes well, would remove completely an existing 10% deficit, but which if it went badly would increase the deficit from 10% to 20%. Suppose the trustees of the fund judged that this was only acceptable if the chance of winning was at least 70% (this means α=8 or more). In that case, the bond weight should be close to 80% even when a future fund surplus is valued.

The key take-away here is that risk appetite is central to investment decisions. But quantifying risk appetite is tricky because the upside and downside risks for trustees are so different. We hope that our way of calibrating attitudes based on valuing α and β is useful here.

Exhibit 10

Effect of Risk Aversion — No Benefit in Having a Future Fund Surplus (β=∞)

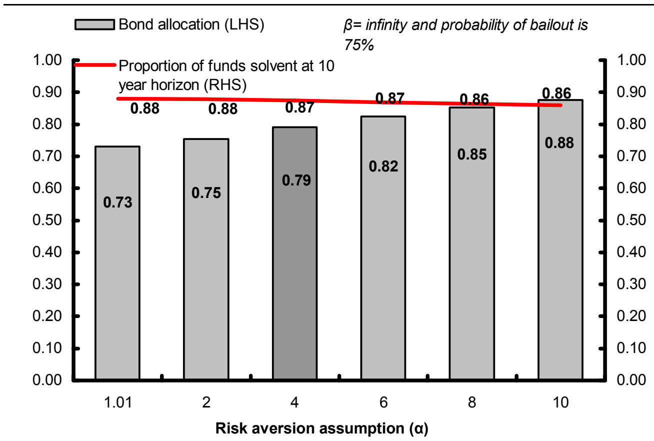

Source: Morgan Stanley Research

Exhibit 11

Effect of Risk Aversion — Benefit Attached to Future Fund Surplus (β=10)

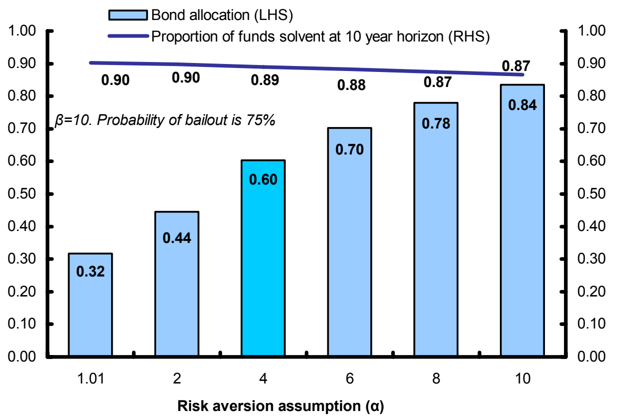

Source: Morgan Stanley Research

3. Longevity risk (see Exhibits 12-13). In the central case we consider three equally likely outcomes for post-retirement life expectancy 10 years ahead: there is a one-third chance of being exactly right so that that people will on average live another 20 years after retiring. But there is a one-third chance that life expectancy is either 20% higher (on average they live for 24 years) or 20% lower (so that on average they only live 16). If the deviations from the average expected outcome are larger, at 30% as opposed to 20%, the optimal bond weight falls from 80% to 75% (when future outperformance beyond 100% funding is not valued) but is little changed when there is a positive benefit to having a future surplus (β=10).

Exhibit 12

Effect of Life Expectancy — No Benefit in Having a Future Fund Surplus (β=∞)

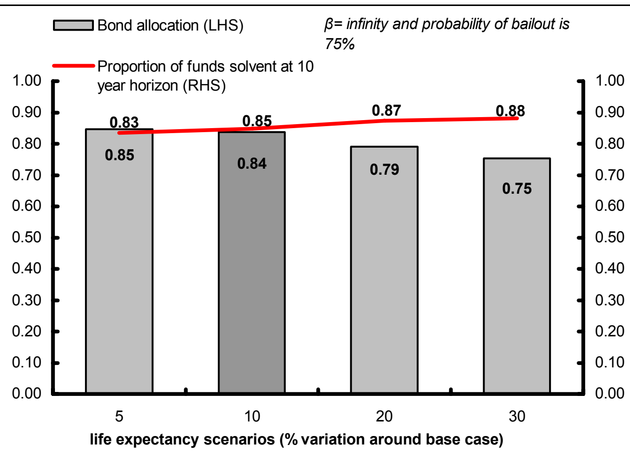

Source: Morgan Stanley Research

Exhibit 13

Effect of Life Expectancy — Benefit Attached to Future Fund Surplus (β=10)

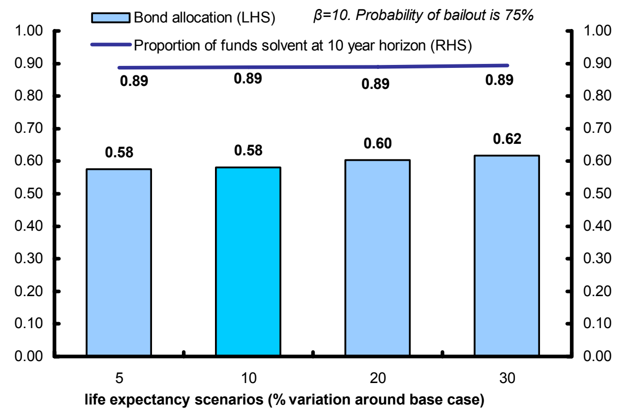

Source: Morgan Stanley Research

The key take-away here is that more longevity risk — which we assume cannot he hedged, and certainly not with bonds — does NOT mean that a more conservative investment strategy is optimal.[12]

4. Bond yields (see Exhibits 14-17). It takes a fairly large shift in yields today to have a dramatic impact on portfolio allocation so long as we assume that the *expected* level of yields T (=10) years ahead is equal to whatever today's yield is. If we keep to the assumption that *on average* 'tomorrow's' yields are the same as today's yields, then at 3% yields today, as opposed to the base case of 1.5%, the optimal bond allocation rises from about 80% to 86% (when outperformance beyond full funding is not valued) and from 60% to 70% when it is valued (β=10).

But there is a more powerful impact on portfolio allocation if higher yields are expected in future while today's yields stay at 1.5%. If today's yield is 1.5% but the expected level of yields in T period's time is 3%, then the optimal bond weight becomes 74% when there is no reward to a surplus and is only 49% if surpluses at the relevant horizon are rewarded (this is against bond weights of almost 80% and 60% respectively if yields are expected, *on average*, to remain at 1.5%). When bond yields are expected to fall, optimal portfolio weights are higher.

The key take-away here is that whether yields are likely to rise (or fall) may be more important than whether they are low (or high) relative to returns on other assets.

[12] This is a result consistent with the results from the report Orszag and others recently presented to the Institute for Actuaries (Orszag et al, 2005).

Exhibit 14
Effect of Bond Yield — No Benefit in Having a Future Fund Surplus (β=∞)

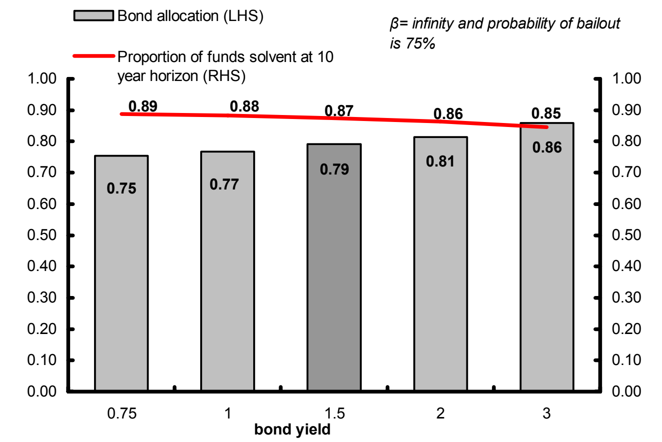

Source: Morgan Stanley Research

Exhibit 15
Effect of Bond Yield — Benefit Attached to Future Fund Surplus (β=10)

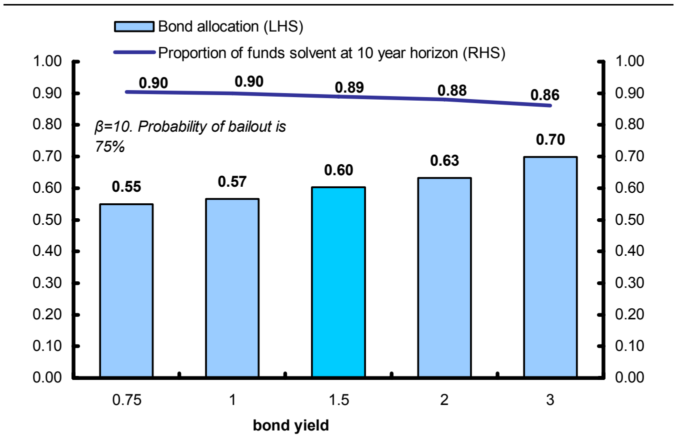

Source: Morgan Stanley Research

Exhibit 16
Effect of Expected Bond Yield (Yield 10 Years Ahead) — No Benefit in Having a Future Fund Surplus (β=∞)

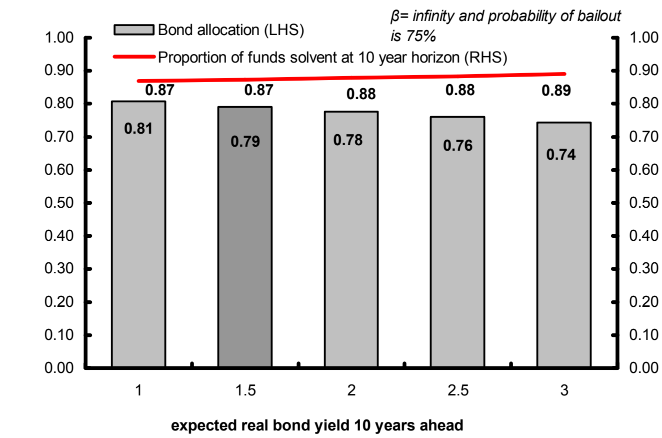

Source: Morgan Stanley Research

Exhibit 17
Effect of Expected Bond Yield (Yield 10 Years Ahead) — Benefit Attached to Future Fund Surplus (β=10)

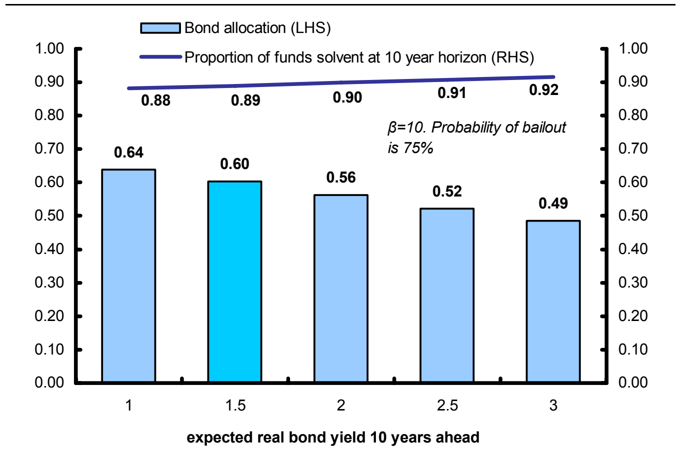

Source: Morgan Stanley Research

5. Level of funding (see Exhibits 18-19). The larger the deficit today (the smaller the value of assets relative to the expected present value of liabilities), the more 'equity' risk is taken. But the impact is not very powerful. At 75% funding, the optimal equity weights (for β infinite and 10 respectively) are 28% and 41% against values of around 20% and 40% at 90% funding. Starting from 100% funding, the optimal weights in equities are 15% and 35% (again versus values of about 20% and 40% at 90% funding).

The key take-away here is that the right investment strategy may not be very sensitive to the current funding level when the appropriate horizon for judging performance is long (say 10 years).

Exhibit 18
Effect of Funding Level — No Benefit in Having a Future Fund Surplus (β=∞)

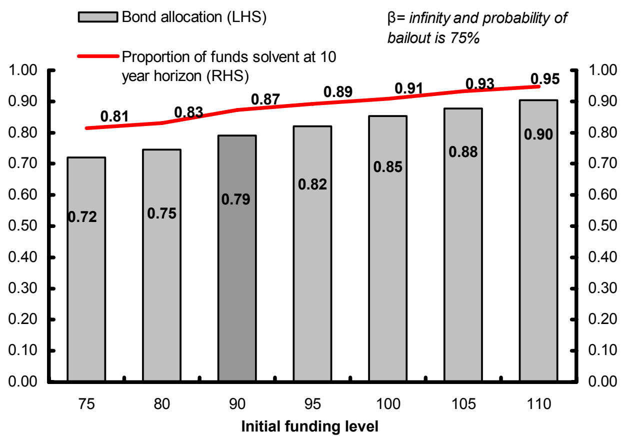

Source: Morgan Stanley Research

Exhibit 19
Effect of Funding Level — Benefit Attached to Future Fund Surplus (β=10)

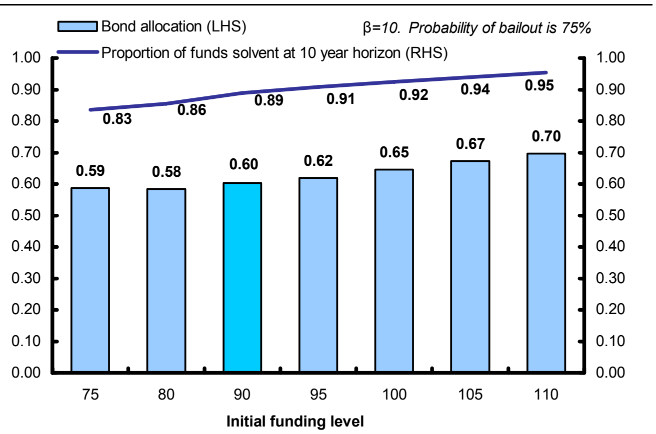

Source: Morgan Stanley Research

Perhaps surprisingly, some factors have very little impact on optimal asset allocation. In particular, the volatility of bond yields, for a given mean expected level of bond yields, has very little effect.

It is surprising and interesting to note that a higher or lower chance of the sponsoring company bailing out the fund has no impact on the right asset choices when the trustees attach literally no weight to the value of outperformance beyond a full funding level (see Exhibit 20). But when weight is attached to such outperformance, the chance of the sponsoring company being willing and able to fund any deficits in the future has a significant impact. A fall in the chance of the sponsoring company being around from 75% to 50% (at β=10) increases the optimal bond weight from 60% to 68%. In the extreme case that there is virtually no chance of a future top-up, the bond weight is higher again at 74%.

Exhibit 20
Effect of Bailout Probability on Asset Allocation

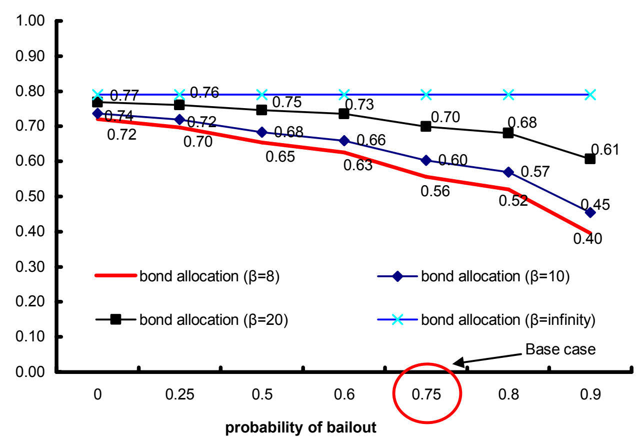

Source: Morgan Stanley Research

One factor that we found to have relatively little impact on the optimal investment strategy is the degree of correlation at longer horizons between bond yield and equity prices. Even when we set the correlation between bond yields and equity returns over a 10-year horizon at -0.6 (which is greater than the degree of correlation between the two based on UK yields and equity returns since 1800), it made only a small difference to investment choice. However, the negative correlation between yields and equity returns does make equities a bit more attractive, since it means that if yields fall (and the cost of pension promises is higher), equities tend to have better-than-average returns.

Conclusions

Our aim has been to think in a systematic way about how variations in the economic environment and in the current structure of pension fund assets and liabilities should influence investment decisions. We have done so using a model that, while allowing for many factors, is nonetheless very abstract; it does not allow for important real world complexities. But we believe that it can be an aid to thinking about the way in which pension schemes should make decisions.

We reach a range of conclusions

The appropriate investment strategy will not be the same across DB pension funds. Those with a strong sponsor, and where trustees attach weight to asset outperformance that moves the fund into surplus at the relevant horizon, will want to hold fewer bonds. Those that attach little weight to higher returns on assets, beyond the point at which the fund is in balance, will prefer much higher bond weights. But despite the likely variability across funds in optimal asset allocation, we find relatively few situations where a bond weighting as low as is currently typical for UK pension funds is optimal.

The combinations of risk aversion and expectations about the future that would justify current low bond weightings look a stretch to us. Currently, a typical UK DB pension fund holds around 30% of asset in bonds, 60% in equities and 10% in alternative assets and cash. We find that such a low weight in bonds is only likely to be optimal for a fund that both expects the sponsoring company to top up the fund in the future with a high probability (85% or more) *and* simultaneously expects to share to a significant extent in any upside that generates surpluses in the future. *In some sense, that combination of characteristics is problematic if, as seems likely, the quid-pro-quo for a sponsor providing a very strong backing to a pension scheme if it faces a deficit is that a disproportionate share of the benefits from asset out-performance accrue to the sponsor.* An alternative combination of expectations about the future and current risk aversion that would justify current low bond weightings also looks a stretch to us. In the case where running a surplus is regarded as a benefit (β=10) and where the probability of bailout is 75%, a bond weighting below 40% could be justified if: a) trustees were not significantly risk averse (α=1), which seems very unlikely; or b) the standard deviation of equity returns was expected to be less than around 15%.

Overall, we find that under a fairly wide range of relevant conditions, a substantial degree of hedging against interest rate risk is likely to be desirable. The simplest interpretation of this result is that long-duration bonds should form a substantial part of pension portfolios. But the same end can be achieved by less direct means — through swaps and other derivatives.

This finding does not reflect pension regulations or accounting rules. The question we pose is what is the right portfolio of assets based on economic factors, and in isolation from any constraints from regulations and accounting. We find that the answer to this question — under many conditions — is that more hedging against interest rate risk should probably be done than is currently typical.

This is significant because many view pressures on trustees and advisors to hold more bond-like assets as pushing them towards decisions that on economic grounds are not very sensible. Our results suggest that this is not generally the case.

The aggregate implications of our analysis are in some sense rather clear — we would expect a fairly substantial switch towards higher bond weightings for most DB pension schemes, particularly in the private sector where the strength of the backing of the sponsor company is likely lower than in public-sector funded schemes (e.g., local authority pension schemes). But the asset price implications are not so clear. Whether such a switch towards more long-duration bonds means that yields will stay at their current low levels — or indeed move lower — is far from obvious. It depends critically on sources of new supply of long-duration bonds. It is perfectly possible — and indeed natural — that greater supply from the public and corporate sector can match the greater demand from re-balancing within pension schemes. **In fact, a switch from equities towards bonds achieved through companies buying back their shares held in other corporate pension funds, and financed through issuing bonds, is neutral. It creates as much new bond supply as is demanded and as much new equity demand as is supplied.** (That argument is spelled out in greater detail in *Real Yields, Pensions and Shifts in Demand for Bonds*, D Miles et al, July 4, 2005.) We remain sceptical that pension fund re-balancing means that real bond yields have to stay at recent, unusually low levels.

Finally, we emphasise again that we have developed an abstract and simplified model. It will be useful to consider generalisations in future work. Allowing investment decisions to be more dynamic — allowing for decisions to be made at regular intervals — is an important generalisation.

Appendix

The full structure of the model and solution technique

We denote the market value of pension fund assets today by A_t. Those funds can be invested in two broad classes of assets: those with returns linked to bonds, and alternative assets — which for simplicity we call 'equities' (Eq_t). We call the asset with returns linked to bond yields 'bonds' (B_t), while recognising that they may involve holding combinations of assets with characteristics not entirely bond-like but which with derivatives can mimic bonds.

Thus: $A_t = B_t + Eq_t$ (1)

B_t = market value of bonds held at time t. This is the number of bonds (n) times a current bond price P_{bt}

Eq_t is the market value of equities held at time t.

The liabilities are commitments to pay pensions in the future. We define a horizon which we think of as the average time to when those pensions are first paid as members retire. We denote this time (measured in years) by T.

At horizon T, a commitment to pay a pension that gives a constant real flow of income has a value that we model as proportional to a level annuity. If life expectancy for a newly retired person T period ahead is M and if the real yield on bonds T periods in the future is y_T then we take the value of an annuity (which we denote by LT, for liability value at time T) to be:

$LT = \Sigma^{M}_{i=0}\ 1/(1+ y_T)^i = (1+ y_T)/ y_T \times [1- 1/(1+ y_T)^M]$ (2)

Both y_T and M are highly uncertain when viewed from time t We are assuming here that the uncertainty over bond yields is about the overall level of yields T periods ahead, and not about the shape of the yield curve at that date. In fact we make the simplifying assumption that the slope of the yield curve is constant and all uncertainty is about the level of yields. (This is why we use a single bond yield at time T, y_T, to discount the cash flows in equation (2)).

The duration of the liability at time T depends upon both y_T and M. If real yields are low, that duration is approximately M/2 (half life expectancy). If it were known in advance that the duration of liabilities at time T would be M/2, then investment in bonds at time t with duration T + M/2 would hedge against most yield curve risk at date T.

So one natural asset to hold at date t is a bond with a duration of approximately T + M/2. Of course, at date t actual life expectancy T periods ahead (i.e., M) is not known. That risk cannot be hedged with bonds. But we assume that there is a zero-coupon bond with a maturity equal to the expected duration of the liabilities. We denote its yield by y_t. The cost of such a bond is naturally

$P_{bt} = 1/(1+yt)^{DUR}$

Where we set DUR to be T + $E_t(M)/2$ (where $E_t(M)$ is the expectation at time t of what life expectancy will be at time T). If the error on life expectancy turns out to be small, this will be roughly the right duration, ex-post (provided the level of real yields is not much more than 3%).

The value of liabilities today, from which we define the funding ratio, is simply the expected present value of LT based on today's bond yields and today's view of expected life expectancy, $E_t(M)$. This is:

$PV (LT) = [((1+ y_t)/ y_t) \times [1 - 1/(1+ y_t)^{Et(M)}]] / (1+ y_t)^T$ (3)

We assume that bond yields at time T are, when viewed from today (time t), uncertain: they are assumed to be equal to yields at time t (y_t) plus a random element and plus a possible factor that reflects anticipated changes in yields.

Thus:

$y_T = y_t + \text{shift} + u$ (4)

Where 'shift' reflects anticipated overall movements in the yield curve and u is a normally distributed random shock with variability σ_u. In the base case we set shift to zero.

The return on equities is also random. The level of stock prices is assumed to be log normal — in other words, logarithmic returns are normally distributed. The level of log equity returns over a horizon of T years (starting from date t) is:

$Re = T \times \mu + \varepsilon$ (5)

μ is the expected *annual* rate of return on equities.

ε is a random (normally distributed) shock with standard deviation of $\sqrt{T} \times \sigma$ (where σ is the *annual* standard deviation of equity returns and we assume the random part of stock returns are not correlated from one year to the next).

The pension fund has to chose B_t and Eq_t, given the current value of assets A_t.

We assume that the goals and interest of the pension fund are described by a utility function which depends on the balance between assets and liabilities at horizon T. Liabilities at horizon T are given by equation (2). Assets are given by:

$$A_T = B_T + Eq_T \qquad (6)$$

The value of each bond held at T (which when it was bought at time t was $1/(1+y_t)^{DUR}$) is $1/(1+y_T)^{E_t(M)/2}$

Where $E_t(M)$ was the expected length of life at time T and $DUR_t = T+ E_t(M)/2$. The maturity (and duration) of the bond bought at time t is therefore $E_t(M)/2$ T periods later. Thus:

$$B_T = n / (1+y_T)^{Et(M)/2}$$

The value of equities at time t is just the value at time t adjusted for the actual rate of return:

$$Eq_T = Eq_t (\exp(T \times \mu + \varepsilon))$$

The value of liabilities at time T depends on life expectancy at that point, which we denote M, and the level of yields at T, y_T. Life expectancy at T can differ substantially from what it was expected to be at time t — so that M and $E_t(M)$ can be very different. Specifically, we assume that actual life expectancy at time T can be substantially above, substantially below or equal to its expectation at time t. The chances of each of these three outcomes are assumed to be equal and the risk around the forecast made at time t is symmetric. Life expectancy could be λ% higher or lower than its forecast value at time t. Thus the three possible outcomes for M are:

$M= (1+ \lambda) \times E_t(M)$ with probability 1/3

$M= E_t(M)$ with probability 1/3

$M= (1 - \lambda) \times E_t(M)$ with probability 1/3

The value of pension liabilities at time T, denoted LT, is:

$$LT = \Sigma^{M}_{i=0}\ 1/(1+ y_T)^i = (1+ y_T)/ y_T \times [1 - 1/(1+ y_T)^M] \qquad (7)$$

The utility function is as follows[13]:

For $A_T < L_T$ utility is

$$(A_T/L_T)^{1-\alpha} / (1-\alpha) \qquad (8)$$

For $A_T > L_T$ utility is

$$(A_T/L_T)^{1-\beta} / (1-\beta) + 1/(1-\alpha) - 1/ (1-\beta) \qquad (9)$$

The extent to which the pain of deficits rises with the shortfall is determined by α. This ties down the extent of risk aversion over downside risks. The degree to which over-funding creates any value at time T is reflected in the value of β. If, as seems highly likely, $\beta > \alpha$, the value of outperformance is much less than the pain of a deficit. If $\alpha > 0$, trustees are risk averse — which would imply that the pain of a shortfall of 2x% is more than double the pain of a shortfall of x%.

Solution technique

We find the value of bonds and value of equities at time t to maximise the expected value of utility at time T. Maximisation of expected utility is standard in finance. It involves calculating the probability-weighted value of utility. We use Monte Carlo methods. We take 2,000 draws from the random components of bond yields and equity prices (u and ε). For each of these we allow for three possible values of longevity (with equal probability), and also two outcomes (sponsor top-up or not) in those cases where E< L. So, given three outcomes for life expectancy and 2,000 outcomes for asset prices we have a total of 6,000 outcomes for the relative values of assets and liabilities at date T. For those of the 6,000 where assets are less than liabilities, we take a probability of p that assets are topped up to the value of liabilities. We can allow for the random shocks to bond yields and equity prices to be correlated.

Bibliography

Haberman, Day, Fogarty, Khorasanee, McWhirter, Nash, Ngwira, Wright and Yakoubov: *A Stochastic Approach to Risk Management and Decision Making in Defined Benefit Pension Schemes*, paper presented to the Institute of Actuaries, January 2003.

Cardinale, Katz, Kumar and Orszag: *Background Risk and Pensions*, paper presented to the Institute of Actuaries, November 2005.

[13] This formulation of the utility function benefited from a helpful discussion with Dr David McCarthy of Imperial College, London.

Disclosure Section

Morgan Stanley & Co. International Limited, authorised and regulated by Financial Services Authority, disseminates in the UK research that it has prepared, and approves solely for the purposes of section 21 of the Financial Services and Markets Act 2000, research which has been prepared by any of its affiliates.

Global Research Conflict Management Policy

This research observes our conflict management policy, available at www.morganstanley.com/institutional/research/conflictpolicies.

Important Disclosures

This report does not provide individually tailored investment advice. It has been prepared without regard to the circumstances and objectives of those who receive it. Morgan Stanley recommends that investors independently evaluate particular investments and strategies, and encourages them to seek a financial adviser's advice. The appropriateness of an investment or strategy will depend on an investor's circumstances and objectives. This report is not an offer to buy or sell any security or to participate in any trading strategy. The value of and income from your investments may vary because of changes in interest rates or foreign exchange rates, securities prices or market indexes, operational or financial conditions of companies or other factors. Past performance is not necessarily a guide to future performance. Estimates of future performance are based on assumptions that may not be realized.

With the exception of information regarding Morgan Stanley, reports prepared by Morgan Stanley research personnel are based on public information. Morgan Stanley makes every effort to use reliable, comprehensive information, but we do not represent that it is accurate or complete. We have no obligation to tell you when opinions or information in this report change apart from when we intend to discontinue research coverage of a company. Facts and views in this report have not been reviewed by, and may not reflect information known to, professionals in other Morgan Stanley business areas, including investment banking personnel.

To our readers in Taiwan: This publication is distributed by Morgan Stanley & Co. International Limited, Taipei Branch; it may not be distributed to or quoted or used by the public media without the express written consent of Morgan Stanley. To our readers in Hong Kong: Information is distributed in Hong Kong by and on behalf of, and is attributable to, Morgan Stanley Dean Witter Asia Limited as part of its regulated activities in Hong Kong; if you have any queries concerning it, contact our Hong Kong sales representatives.

This publication is disseminated in Japan by Morgan Stanley Japan Securities Co., Ltd; in Canada by Morgan Stanley Canada Limited, which has approved of and takes responsibility for its contents in Canada; in Germany by Morgan Stanley Bank AG, Frankfurt am Main, regulated by Bundesanstalt fuer Finanzdienstleistungsaufsicht (BaFin);in Spain by Morgan Stanley, S.V., S.A., a Morgan Stanley group company, supervised by the Spanish Securities Markets Commission(CNMV), which states that it is written and distributed in accordance with rules of conduct for financial research under Spanish regulations; in the US by Morgan Stanley & Co. Incorporated and Morgan Stanley DW Inc., which accept responsibility for its contents. Morgan Stanley & Co. International Limited, authorized and regulated by Financial Services Authority, disseminates in the UK research it has prepared, and approves solely for purposes of section 21 of the Financial Services and Markets Act 2000, research prepared by any affiliates. Private UK investors should obtain the advice of their Morgan Stanley & Co. International Limited representative about the investments concerned. In Australia, this report and any access to it is intended only for "wholesale clients" within the meaning of the Australian Corporations Act.

MORGAN STANLEY RESEARCH

The Americas
1585 Broadway
New York, NY 10036-8293
United States
Tel: +1 (1) 212 761 4000

Europe
25 Cabot Square, Canary Wharf
London E14 4QA
United Kingdom
Tel: +44 (0) 20 7 425 8000

Japan
4-20-3 Ebisu, Shibuya-ku
Tokyo 150-6008
Japan
Tel: +81 (0) 3 5424 5000

Asia/Pacific
Three Exchange Square
Central
Hong Kong
Tel: +852 2848 5200