



DISCUSSION PAPER PI-0809

By how much can a Diversified Approach to Investing Improve the Prospects of Reducing a DB Pensions Deficit?

Andrew Brigden, Andrew Clare and Shamik Dhar

April 2008

ISSN 1367-580X

The Pensions Institute
Cass Business School
City University
106 Bunhill Row London
EC1Y 8TZ
UNITED KINGDOM

<http://www.pensions-institute.org/>

By how much can a diversified approach to investing improve the prospects of reducing a DB pension deficit ?

Andrew Brigden,¹

Andrew Clare

&

Shamik Dhar

Abstract

Following the decline in global equity markets, the rise in bond prices and the downward revisions to assumed mortality rates between 2001 and 2003, the UK's defined benefit (DB) pensions industry went from a situation where surpluses and scheme sponsor contribution holidays were commonplace, to a situation where fund deficits were the norm. This is the situation that persists today. In this paper, we present a model of a typical DB pension scheme, where we take explicit account of the linkages between both sides of the pension fund balance sheet. We use this model to assess the likely benefits of taking a more diversified approach to asset allocation compared with the traditional heavy reliance on UK equities. Our results show that there are clear gains to be had from adopting a more diversified approach to pension fund asset allocation in terms of the reduction in scheme funding volatility over time.

Keywords/phrases: Alternative asset classes; diversification benefits; defined benefit pension scheme; pension fund deficit; asset allocation

JEL Codes: G10, G11 and G23

¹Authors' affiliations: Andrew Brigden (andrew.brigden@fathom-consulting.com) and Shamik Dhar (shamik.dhar@fathom-consulting.com), Fathom Financial Consulting, 48, Gresham Street, London. EC2V 7AY; and Andrew Clare (a.clare@city.ac.uk), Cass Business School, 106, Bunhill Row, London. EC1Y 8TZ .

1. Introduction

The UK's defined benefit pension industry remains a crucial part of private sector pension provision in the UK. But between 2001 and 2003/4 the industry became caught up in what has been described as a "perfect storm" (see Ritchie and Dumbreck (2007)). This "perfect storm" comprised a sharp decline in global equity prices between 2001 and 2003, rapidly rising bond prices and a reassessment of mortality assumptions used to project future pension payments.

The shock to equity prices had a large impact because the majority of pension fund assets were held in this asset class. At the end of 2001 equities made up around 71% of total long-term asset holdings (UBS, 2006). Meanwhile, rising bond prices, delivered a 'double whammy'. First, new UK accounting standards that came into force over this period (FRS17) recommended that liabilities be valued by discounting future pension payments by using the prevailing yield on AA-rated sterling corporate bonds. Because of this move to marking to market pension liabilities the subsequent fall in sterling corporate AA yields led to an increase in the discounted present value of liabilities and a further widening of fund deficits. Second, although falling AA yields led to an increase in the value of any sterling corporate bonds that formed part of pension portfolios, those corporate bonds held by pension funds had average Macaulay duration of around 7 to 8 years. In contrast, the Macaulay duration of typical scheme liabilities was often well in excess of twenty years. This duration mismatch meant that the rate of increase of the present value of scheme liabilities was much faster than the rise in the value of bond holdings. Some schemes tried to remedy the duration mismatch by buying much longer dated bonds – both corporate and UK government bonds – but these bonds were and still are in relatively short supply, and the increase in demand for them pushed prices higher and yields lower, thus causing liability values to rise further still, in what can be viewed as a largely self-defeating exercise. Finally, although it is questionable whether mortality trends change as rapidly as financial markets, the soaring bond and collapsing equity prices focussed actuarial attention on the longevity assumptions that they had been relying on up until this point. Uncomfortable with the generous assumptions that had been made in the past, scheme actuaries moved to more conservative assessments of mortality rates, revising them down sharply for many schemes. This had the effect of raising the value of estimated scheme liabilities so much so that the move effectively offset much of the beneficial impact of the recovery in global equity markets since March 2003 for many schemes.

Widespread scheme deficits are commonplace today within the UK's DB industry. According to a survey conducted by the UK's Pensions Regulator of 5,800 DB schemes (reckoned to

be around half the total number in existence) 83% of the schemes were in deficit in March 2006 (Pensions Regulator (2006)). The same survey showed that 58% of these schemes were reported to be closed to new members, with around 1,150 of these schemes closing between 2001 and 2003. The closure of a scheme of course exacerbates the funding issue since, by definition, the scheme will mature more rapidly than if it were to remain open to new, younger employees.

Given all this, many trustee boards, scheme sponsors and their advisors are looking to identify ways to return their schemes to fully funded positions as soon as possible, while simultaneously seeking to reduce the cost and risk inherent in any solution. To achieve these competing ends it is necessary to consider all aspects of the DB funding issue within one coherent framework. These aspects would naturally include both scheme assets and liabilities, the correlations between the two sides of the balance sheet and the contribution rates of scheme members and of the scheme sponsor. Blake (2003) argues that the Myners report gave the pensions industry the initial spur to consider the correlation both between asset classes and between asset classes and scheme liabilities over time. In this paper we present and utilise a model that takes account of both sides of the pension fund balance sheet, with the specific aim of identifying the likely benefit in terms of risk reduction of holding a more diversified set of asset classes.

Our results indicate that even given quite conservative assumptions about future asset class returns and risks, adopting a more diversified approach to pension fund asset allocation that could lead to a considerable reduction in the volatility of scheme funding. We estimate a significant drop in the probability of a scheme being underfunded in the future by moving from a traditional 60/40 equity/bond allocation to one which encompasses between 30% to 50% of alternative asset classes, where we define alternatives as including an allocation to high yield, commercial property, commodity futures and to two hedge fund categories – market neutral and macro.

The rest of this paper is organised as follows: in section 2 we briefly review the related literature; in Section 3 we describe the model that we use to define a typical DB scheme; in section 4 we describe the asset class data used in the study; in section 5 we present the results of our monte carlo-based experiments; and finally, in section 6, we draw conclusions and suggest areas for future research.

2. Modelling the DB pension problem

Despite the apparent severity of the funding situation faced by many of the UK's DB schemes today, there has been relatively little academic work in this area aimed at modelling the problem. Of related interest to the DB problem is an exercise undertaken by Blake *et al* (2001) who develop an analytical model of a Defined Contribution (DC) pension plan, where the "benchmark" was the retirement income that could be derived from membership of a typical company DB plan. Using stochastic simulations that take account of the various, relevant key risk factors such as asset returns and interest rates, but also the risks associated with an individual's future earnings and the risk of unemployment, the authors experiment with alternative asset allocation rules which are designed, *ex ante*, to achieve a certain, targeted retirement income. Using relatively conventional asset classes, they find that their results are highly sensitive to the choice of asset allocation strategy. Furthermore, Blake *et al* (2001) concluded – rather prophetically – that the risk inherent in a DC plan for an individual was far greater compared with an "equivalent" DB plan. That is, given an equivalent length of service in the scheme and similar member contribution rates, the risks around achieving a particular level of pension was found to be greater in the case of the DC option. This is because the plan sponsor bears nearly all of the risks in the DB pension plan; risks that all materialised for DB plan sponsors in the early part of this century. Haberman *et al* (2003) develop a model of a defined benefit scheme to investigate issues relating to contribution rates, asset allocation decisions and which allows for the development of risk-based performance measures and analytics. Finally, using a model of a representative open DB scheme Driver and Selvaggi (2007) explored various issues relating to: the choice of the discount factor used to evaluate the present value of liabilities; the possible use of cash flow analysis as an alternative to the current mark to market approach to analysing the DB funding problem; and the sensitivity of the funding position to changes in pension benefits.

In this paper we describe and use a cohort-based model of a typical closed DB pension scheme² to analyse the potential benefits that might accrue to both scheme members and sponsors from adopting a more diversified approach to the management of pension fund assets as trustees seek to reduce their pension deficits. In particular we will consider how the combination of traditional and alternative asset classes such as hedge funds and commodities might improve the risk return trade off. Other researchers have investigated the possible benefits of diversifying portfolios in this manner, for example, see Amenc *et al* (2002) or Schneeweis *et al* (2002). These studies tend to show that the addition of alternative asset classes to a portfolio containing more traditional asset classes, leads to an

²Although the model can be adapted so that it represents an open scheme too.

expansion of the mean-variance frontier. But as far as we are aware no academic work to date has investigated the possible benefits of this diversification within the rich DB framework that we develop and use here, and which also takes explicit account of the correlation between asset class returns and the changing value of scheme liabilities over time.

The model used here differs from the one constructed by Haberman *et al* (2003). The approach taken by these authors is essentially analytical. Their approach lends itself easily to stochastic analysis. By contrast, the model used in this paper to analyse similar issues is based upon a “bottom-up” aggregation of the pension rights of both past and present scheme members that are all represented by cohorts. This “bottom up” approach mimics closely the way in which the liabilities of a typical scheme are constructed in practice by consulting actuaries. Our approach potentially allows for greater fine tuning of certain parameters, allowing us, for example, to change the earnings or mortality profiles for particular cohorts enabling us to capture a very diverse range of cohort characteristics which, taken together, can facilitate the close approximation of the main features of any real world scheme.

The model described in Section 3 below is in fact a direct descendent of the model used by the Driver and Selvaggi (2007) in their research, and which was developed for the Association of British Insurers (ABI) by the authors of this paper. However, our main purpose here is to explore the relationship between asset allocation and key DB funding metrics, an aspect of the funding issue that was not explored by Driver and Selvaggi (2007). In addition, the model presented in this paper differs from the one used in the ABI’s work in that we incorporate stochastic modelling techniques to capture the return and risk characteristics of a typical DB scheme’s assets and their relationship with the scheme liabilities going forward. By introducing these stochastic elements to the model we can use it to focus on the risks around projected funding positions from various asset allocation starting points.

3. A cohort-based model of a defined benefit pension plan

We model a typical Defined Benefit (DB) pension scheme. These pension arrangements are also known as final salary schemes, since the member’s pension is a direct (and defined) function of their final salary³. The purpose of this section is to describe this model and to outline the features of a typical scheme which will act as a base case for the asset allocation simulations that are described in Section 4

³ See Blake (2006) for a description of typical pension scheme arrangements.

We begin by defining a set of N representative scheme members. There is a single scheme sponsor and all new employees are enrolled into the scheme for as long as the scheme remains open. Each member is representative of a particular scheme cohort and, through the manipulation of scaling factors, can represent just one scheme member or a group of members that all have the same basic characteristics. For example, one cohort might consist of women members aged 32 at their next birthday, while another may represent a male members that are aged 45 at their next birthday etc. By scaling up the number of people within each cohort we can potentially reflect the structure of a wide range of actual schemes. Each scheme member (or cohort) in the model accrues pension rights according to the following expression:

$$P_{i,r} = S_i \times \frac{n_i}{A} \quad (1)$$

where P_i represents the gross annual pension for the i th scheme member, to which the member is entitled on reaching the pre-defined retirement age, r , for the rest of their lives. S_i represents the final salary of the i th scheme member. n_i represents the number of full years that the i th scheme member has contributed to and been a member of the scheme. And A represents the scheme accrual rate, for example, 60, 80, 100 etc. The firm's pensionable age is distinct for males and females, and can be varied through time. For our base case we set the retirement age for women at 60 and for men at 65.

For experiments that follow in this paper, we need to define an age and gender structure for the staff in the scheme, both retrospectively and prospectively. For our benchmark, or base case scheme the number of staff within the scheme remains constant over time and is benchmarked to an estimate of the age and gender structure of the UK workforce as a whole during the period June to August 2006 reported by the ONS in the October 2006 Labour Market Statistics First Release. In this publication, the ONS provide estimates of the number of people in employment who fall into one of six age ranges. For example, the First Release contains estimates of the number of people in employment who are between 16 and 17, between 18 and 24, between 25 and 34 and so on. Separate estimates of the number of males and females in employment by broad age group are also available. In order to construct estimates of the number of males and females in employment according to age at last birthday, running from 16 through to 64, a polynomial was fitted through the ONS data. This estimated age and gender structures of the scheme based upon this ONS report are shown in Figure 1.

The representative scheme sponsor faces an exogenous quit rate, q . Each period, a proportion, q , of each age cohort within the firm, whether male or female, resign to take up employment elsewhere. For the benchmark scheme we set q equal to 0.15, which means that every year 15% of the workforce quit to take up employment elsewhere (or are made redundant). At this point, these former employees become deferred members of the scheme. Within the model these individuals are replaced by a person of the same age and gender. Independently of these resignations, a number of men and women above a defined age threshold leave the firm and do not seek further employment. Effectively, these people are assumed to have chosen early retirement; they have stopped working before reaching the age at which they are entitled to a pension. The age and the rate at which people take early retirement can be inferred from the age/employment profiles in Figure 1. The number of males in the firm peaks at age 48. From the age of 49, males start to leave the firm at a rate given by the slope of the male age/employment profile until we reach a point where there are no longer any male members of staff. This occurs at age 62. The number of females in the firm peak at age 47. From the age of 48, females start to leave the firm at a rate given by the slope of the female age/employment profile until we reach a point where there are no longer any female members of staff. This occurs at age 60. Taken together, these two assumptions, that those leaving to take up employment elsewhere are immediately replaced by a person of the same age and gender, and that people retire at an age and rate given by the downward sloping portion of the age/employment profiles, ensure that the age and gender structure of the firm is preserved⁴.

In our base case we assume that the scheme opened in 1970 and closed in 2003. However the model is flexible enough to allow us to vary these dates, and to model a scheme that remains open. While the scheme is open, staff joining the firm are automatically enrolled in to it. Once the scheme is closed, existing staff remain in the scheme until they leave the firm, but new staff joining the firm are not enrolled in it. This means that the exogenous quit rate is an important determinant of the magnitude of scheme liabilities. If the quit rate is high then once the scheme has closed the number of active members will drop off more quickly than if the quit rate is low.

We assume that the age/earnings profiles for both male and female staff within the firm were the same as the national averages recorded in the April 2005 Annual Survey of Hours and Earnings (ASHE). These age/earnings profiles are shown in Figure 2. The ASHE provides information on the distribution of earnings of males and females within seven broad age

⁴These are all parameters that can be varied depending upon the need to replicate different age and gender structures.

groups. In order to construct estimates of earnings of both males and females according to age at last birthday, running from 16 through to 64, we fitted a polynomial through the ASHE data points. A simple shift parameter allows us to move these age/earnings profiles up or down according to whether the firm we are modelling pays above or below the national average salary. It should also be noted that according to the ASHE survey average earnings begin to fall for men at age 42 and for women at age 33. This is because this is an economy wide average and reflects the fact that some older high income individuals choose to retire early, therefore bringing the average earnings down. We believe that for most workers continuing in full time employment within a DB scheme this earnings profile is probably not appropriate. We therefore “flatten off” the profile at its peak for both men and women (shown by the dotted lines in Figure 2). This means that there are no more age-related salary increases for our cohorts once they reach the peak of their earnings in terms of age, but also that there is no age related decline in earnings either.

Because our representative scheme began life in 1970 we need to construct a salary history for scheme members, and a projection of future salaries for current scheme members, even though the scheme is closed to new entrants from 2003 onwards. We determine the past, current and future salaries for each cohort in the following way. For earnings before 2005, we deflate the 2005 ASHE data by an assumed trend rate of productivity growth of 2.25%, and by the observed rate of RPI inflation. This gives us the historic earnings profile of past and current scheme members. For earnings beyond 2005 we inflate the 2005 ASHE data by the same assumed trend rate of productivity growth, and by an assumed rate of future RPI inflation (2.5%). This gives us the future earnings profiles of current scheme members if the scheme is closed, or of current and future scheme members if the scheme is open. This technique effectively preserves the shape of the age/earnings profile through time.

The pension accrual formula shown in expression (1) can be augmented to allow for a variety of index-linking. Some UK pension schemes that had accumulated surpluses during the 1990s linked all future pensions in payment to the RPI. New pension legislation in 1995 made it compulsory from April 1997 onwards that all future pension accruals should produce a pension in payment that was linked to RPI with an annual ceiling of 5.0% and a floor of 0.0%. From 2005 new accruals were to be accumulated linking pensions in payment to inflation to a maximum uplift of 2.5% and again with a floor of zero. This indexation is known as limited pension indexation (LPI). Although the scheme member is exposed to positive inflation between the time they leave the scheme and the time they draw their pension, such pensions in payment at least have a degree of inflation protection. In our scheme we “offer” each cohort LPI(0,5) related pension in payment protection, but only for those pension rights

accrued after 1997. This means that a person accruing pension rights before and after 1997 will have part of their final pension inflation proofed up to a limit of 5.0% inflation, and with a promise that the nominal pension will not fall from year to year, given the floor of 0.0%; while the remainder of the pension will consist of a fixed nominal payment.

Some portion of the pension in payment will therefore be uplifted by the retail price index (RPI), from the point at which the member begins to draw their pension:

$$P_{it} = \left(S_i \times \frac{n_i}{A} \right) \times \left(\frac{RPI_t - RPI_{rt}}{RPI_{rt}} \right)$$

$$\text{s.t. } 0.0\% \leq \left(\frac{RPI_{t+1} - RPI_t}{RPI_t} \right) \leq 5.0\%, \text{ or } 2.5\%, \text{ for all } t > rt \quad (2)$$

where P_{it} is the annual pension of member, or cohort i in year t , and the subscript rt represents the first year in which the member, or cohort, can draw their pension.

We have now described the evolution through time of the workforce, of exits from the workforce, and of rates of pay, which vary by age and by gender. We have also outlined the nature of the defined benefit scheme, and the extent to which entitlements are indexed, or partly indexed. All that remains in order to construct projections for the flow of scheme liabilities is to estimate mortality rates applicable to active members, deferred members and to current pensioners. These mortality rates were obtained from two distinct sources. For data up to and including 2004, we used ONS mortality statistics. For data beyond 2004, we used the latest GAD principal projections published in October 2005. The ONS and GAD data were used to construct two matrices, one for males and one for females, containing the probability that an individual born in year i dies in year j , conditional on that individual being alive in year $j-1$. For completeness, we allow i to run from 1900 to 2099 and j to run from 1900 to 2099. Neither the historical ONS data, nor the GAD projections are available for each year of birth, or for each year of death continuously from 1900 to 2099. In other words, we are not able to fill in every cell in our two matrices using the official data. Consequently, we were forced to interpolate, both horizontally and vertically in the two matrices in order to populate each and every cell. For simplicity, we assume that the death probability remains constant, at around 30% per annum, beyond age 100. We also assume that mortality rates for each age group were fixed before 1976. These data allow us to project the length of time that a scheme member can expect to receive their pension in retirement and therefore, when combined with the age, gender, earnings and pension accrual assumptions allows us to calculate the liability that each cohort represents for the scheme sponsor. The aggregation

of these cohort-based liabilities gives the total liabilities which arise from scheme membership. We calculate the present value of the sum of these future liabilities using the approach adopted in current accounting practice and described in FRS17. That is, we use the current market yield on a AA-rated, sterling denominated corporate bond to discount all future liabilities. The sum of each of these discounted obligations is therefore the total present value of the scheme's liabilities at each point in time. This is shown in expression (3):

$$L_t = \sum_{i=1}^{J-i} \frac{E(P_t)}{(1+r_{l,t})^{t-i}} \tag{3}$$

where L_t represents the total present value of liabilities in year t , $E(P)$ is the expected future value of pension payments in year t , r_l is the discount rate applied to all liabilities and i is 2006. J represents the last year in which pensions are due to be paid in the case of a closed scheme.

We assume that these liabilities are met over time by a combination of three sources of revenue. The first two are the contributions of the scheme members and of the scheme sponsor. We assume a contribution rate for members and for the scheme sponsor which is expressed as a proportion of the total salary bill at the beginning of each year. For our base case we set this at 6.0% for both members and for the scheme sponsor. In total then the scheme has received and will continue to receive in the future an injection at the start of each year of 12% of the start-of-year total company salary.

The other source of revenue is derived from the assets held by the scheme. In our base case we assume that the current market value of the assets held by the firm represent only 80% of the total present value of total scheme liabilities. This funding ratio is not untypical today. Over time the value of these assets evolves from year to year according to the following expression:

$$A_t = (1+r_{a,t})A_{t-1} + C_t - P_t \tag{4}$$

where A_t represents the value of scheme assets at the end of year t , $r_{a,t}$ is the rate of return achieved on scheme assets over year t , C_t is the sum of the scheme member and sponsor's contributions to the scheme in year t and P_t represents the pension payments that are made to scheme members in retirement in year t .

In Figure 3 we present the liabilities as faced by our representative pension scheme at the start of the experiment. The chart shows the future values of the liabilities as at 2006, broken down into the three broad categories. In present value terms the largest liability is represented by deferred scheme members who represent 53% of the total. Active members and current pensioners represent 20% and 27% respectively. The total present value of scheme liabilities is just under £100m calculated on a FRS17 basis. The modified duration of the scheme's liabilities is 18%, which is similar to industry estimates of average duration (the scheme's Macaulay duration is 19 years).

4. Asset class data

Expression (3) shows that one of the key components of the scheme's funding position is the rate of return that it achieves on its assets, r_a , and also the volatility of this return over time too. It is the nature of this rate of return that is the focus of this paper. In our analysis we specify a wide range of potential asset classes both traditional and 'alternative' that the scheme can hold and which will therefore affect the nature of r_a . As proxies for the range of asset classes that the hypothetical pension fund can hold we investigate a set of benchmark indices that it is hoped are representative of each hypothesised asset class.

The traditional asset classes consist of cash, gilts, UK and overseas equities, and sterling denominated corporate bonds. The proxy for the rate of return on cash is the one month sterling LIBOR rate. The gilt portfolio is proxied by the total return on the Datastream ten year gilt benchmark index, while the total return data for UK equities was proxied by the FTSE 100 equity index. We also include AA-rated sterling corporate bond index in our set of asset classes, as defined by the relevant total return index as calculated by IBOXX. For the overseas equity portfolio we use three equity indices which is comprised of 60% US equities, 20% European equities (ex UK), and 20% Japanese equities, where we analyse the total return indices calculated by Datastream for each of the three geographical areas. The proportions reflect the sort of exposure that pension funds typically seek in their developed economy equity portfolios.

We also include a range of "alternative asset classes". These include: high yield corporate bonds; UK commercial property; commodity futures; and two hedge fund asset classes – market neutral and macro, as defined by CSFB Tremont. In total there are ten broad asset classes included in our study that range from low risk and traditional (cash and gilts) to higher risk and less traditional (commodities and hedge funds). In this sense our analysis embodies the whole spectrum of asset classes that are currently being considered as

suitable by pension fund trustees. The total return indices that we analyse and that we use to proxy for these asset classes are described in Table 1, along with the indices for the traditional asset classes too.

We run the model described in section 3 over time and monitor the funding ratio of the scheme. Because our focus is on the impact of asset allocation on the hypothetical scheme's funding over time, we therefore need to specify the expected returns on ten asset classes and the components of a ten by ten variance-covariance matrix. There is no 'correct' way to do this. In practice, life companies and other institutional investors with multi-asset class portfolios use a mixture of forward and backward looking analysis to derive the inputs for any mean-variance analysis exercise. The aim of this paper is not to derive an optimal asset portfolio, but rather to show how greater asset class diversity changes key DB funding parameters. Nevertheless, we still need to adopt some method for projecting future returns and for projecting the risks around these expected returns. To this end, we also use a mixture of forward and backward looking analysis.

For the expected returns vector we use as a benchmark the consensus view with regard to the likely expected return on UK cash, gilts, sterling corporate bonds and equities. For UK cash we use an expected return of 5.0%, which reflects estimates of the long term trend rate of growth in the UK economy, and an expectation of long-term inflation. Given the current monetary policy and growth environment, a figure of anywhere between 4.75% and 5.00% would seem to us to be a sensible proxy for the UK's neutral policy rate, and therefore for the expected return on cash. For a portfolio of medium-dated gilts we use an expected return of 5.25%, which implies that the risk premium, largely related to the uncertainty surrounding future inflation, is 0.25%. For AA-rated sterling denominated corporate bonds we assume a spread over gilts of 0.75% per annum. The actual spread over gilts for this asset class since 1998 has been just over 0.80% and at the time of writing was approximately 0.50%. Finally, we assume an equity market risk premium of 2.75%, which yields an expected return on UK equities of 8.00%. Again this is in line with current professional practice and, for example, FSA guidelines on how to represent to investors the likely returns available from products based upon UK equities. The expected returns that we use for these asset classes are summarised in column 1 of Table 2. For the other asset classes presented in Table 2, we base the expected returns around these more conventional asset classes and upon their return history since 1998. Table 2 also discloses the correlation matrix for the ten asset classes. We used data from January 1998 to April 2007 to calculate asset class variances and covariances. We chose this sample period since it embodied at least one business

cycle and therefore will hopefully represent reasonable estimate of volatilities and covariances for the future too.

5. Results

Using the pension fund base case described in Section 3 and with the liabilities in 2006 shown in Figure 3, we hypothesise that the scheme holds the ten asset classes described in Section 4 in a range of proportions with varying amounts devoted to the set of five alternative asset classes. These initial holdings are presented in Table 3. The least adventurous asset allocation is in column two of Table 2. Here gilts and cash comprise 35% and 5% respectively of holdings while the remainder is allocated to UK equities. The most adventurous asset allocation is entitled the “Yale” model, although it is in no way meant to represent Yale University’s investment strategy. But in this asset allocation, 10% is invested in each of the ‘alternative’ asset classes.

For each of these asset allocations we use monte carlo techniques to simulate the experience of the representative pension fund, where we use the expected returns and variance/covariance matrix estimated using the index data, to simulate the returns on these asset classes over time. The returns are simulated under the assumption that they are jointly normally distributed and related of course, via their covariances over time. For each of the asset allocations we use 10,000 replications of asset class returns. The present value of the liabilities also evolve stochastically, since the AA corporate bond discount rate can be recovered from the asset class variance/covariance matrix. We therefore allow financial market conditions to have an impact on the scheme’s funding position over time through its influence upon both r_f and on r_a in expressions (2) and (3) respectively.

In Figure 4 we have presented the results of these simulations for each of the seven asset allocations. The brown line in each of the charts shows the present discounted value of the scheme’s total liabilities as we move through time. The dark blue line represents the expected present value of scheme assets over time. However, we impose on the experiment the following restriction, that is, in any year when the present value of scheme assets reaches 100% of the present value of scheme liabilities, or above, we assume that the scheme switches to a matching solution, and at that point the present value of their assets no longer comprise the average asset value represented by the dark blue line. The dark blue line itself then represents at all points in time the average present value of scheme assets of all those schemes that are still underfunded. We could have chosen a higher proportion, such as 125% at which point the scheme might opt for a full buyout, but for the purposes of

this experiment this assumption does not change the marginal impact of the move to a more diversified investment portfolio.

The blue dotted lines in the charts represent the 90% confidence interval around average scheme values. The upper confidence interval is bounded to be equal to the present value of scheme liabilities since we impose the switching condition at this point. There is no such restriction on the lower bound. The difference between the dotted lines effectively represent the corridor of uncertainty for the scheme's asset values. As the asset allocation gradually becomes more diverse this corridor becomes progressively narrower. We can see this most clearly by comparing the "base case" with the "Yale model" allocations.

The impact of moving to a more diversified approach to asset class investment is perhaps clearer when we consider Figure 5. In this Figure we present the proportion of schemes that remain underfunded for the range of asset allocations. To begin with of course all schemes are underfunded by construction. But by 2027 44% of schemes relying on the 'traditional' asset allocation remain underfunded, compared with 25% with the "Yale model" allocation. However, it is clear that at the start of the experiment the average number of underfunded funds does not vary considerably. The reason for this is that the concentration on developed economy equities produces greater asset return volatility which means that there is a greater chance of getting to a fully funded position, since the spread of possible outcomes is greater. For a more diversified portfolio the volatility is lower and the progression towards full funding is slower, but more certain. This explains why the average number of underfunded schemes is eventually greater in the case of the less diversified asset portfolio ... for every scheme that got lucky by sticking with a high proportion of developed economy equities, there is another that experienced the other side of equity market volatility, that never recovered to get to a fully funded position. This number for the "traditional" asset allocation is around 30%, compared with 8% for the "Yale model" allocation.

The consequences of this "go for broke" strategy, whereby schemes stick with a more volatile asset portfolio, is perhaps better understood by looking at average funding ratios over time. These are shown for the range of asset allocations in Figure 6. By 2017 the average funding ratio is over 95% for the "Yale model" asset allocation, compared with just over 86% for the "Traditional" asset allocation. These averages improve as we increase the diversity of the assets within the portfolio. This metric reaches a "steady state" for all asset allocation strategies eventually. Effectively this steady state is an average of schemes that are fully funded (funding ratio is 100%), and of those that have run out of assets (funding ratio is 0.0%). Under the Yale model approach average funding settles down to 92%. With the

traditional approach the average funding ratio settles down to just under 71%. For the pensions industry as a whole then failing to diversify adequately could lead to a large number of scheme failures (entering into the Pension Protection Fund), significant increases in sponsor contributions, or reductions in scheme member benefits.

Our results are derived without adjustments over time to the starting asset allocation, or to contribution rates etc. In practice higher contributions might be made (if these could be afforded), or scheme benefits might be cut if the scheme experienced a significant fall in the value of its assets in any one year. But the basic results still apply: a more diversified approach to asset allocation other things being equal, could lead to a sizeable reduction in the likelihood that a DB scheme will require additional contributions or will need to cut scheme benefits over the longer term. These results are robust so long as the variance-covariance matrix used in the experiments turn out to be a close approximation to actual variances and covariances going forward.

6. Summary and suggested areas for future research

The approach to the management of Defined Benefit pension schemes has changed quite dramatically over the last few years. During the 1980s and 1990s the objective seemed to be one of return maximisation, which explains why the typical UK DB scheme had a significant reliance on publicly traded equities. However, a whole range of factors combined at the beginning of this century that changed this focus. Today the focus is on return maximisation, with the added goal of risk minimisation. This approach is often referred to as Liability Driven Investment (LDI), and involves careful consideration of the nature of scheme liabilities along with an equally careful consideration of the sort of asset classes that might best meet the pension needs of scheme members, with the lowest probability that these promises will not be met. Large schemes have made use of derivative contracts like swaps to try to reduce the funding risks. However, in this paper we have considered the role that asset allocation can play in helping to reduce this risk.

A more diverse approach to asset allocation and the consideration of alternative asset classes could help to reduce the volatility of total scheme assets over time, as long as the correlations between asset classes are relatively low and ideally if the correlation between asset classes and scheme liabilities is also low, and ideally negative. Our results put some hard numbers on the possible benefits from increased diversification, although they will only be as robust as the expected return and correlation assumptions used in the monte carlo experiments.

We intend to pursue this area of research by developing the model presented here to address other key aspects of the pension fund problem.

The model presented here treats the contribution of the scheme sponsor as exogenous and without risk. In reality the strength of the employer covenant is an important factor in asset allocation decisions, and could be modelled here by specifying a simple stochastic model of a generic firm. Alternatively one could treat the sponsor contributions as a bond, where the discount rate on that bond would be higher the lower the credit quality of the firm. Thus higher discount rates would lead, other things equal, to a fall in the present value of future sponsor contributions and a decrease in the funding ratio.

The asset allocation here is assumed not to change, until the funding ratio equals or exceeds 100%. The framework developed here could be used to explore the role and usefulness of a variety of dynamic asset allocation strategies along with more dynamic funding strategies. The interaction between these could be used to derive the optimal asset allocation for a wide variety of schemes.

An important element of the current regulatory environment for the DB pensions industry is the role of the pension provider of last resort – the Pension Protection Fund (PPF). By specifying a model of the UK economy and combining that with key parameters in the cohort-based model here such as inflation, productivity and interest rates etc, one could use this framework to explore the likely burdens on the PPF going forward, and to test various regulatory rules, where the aim would be to establish rules that limit the potential burden on the PPF without materially reducing the protection that it offers.

Finally, a number of large pension schemes have made use of derivatives as part of their management of the risks that they face. The highest profile example of this being the use that Boots plc made of swaps before the collapse in global equity markets in 2001. The framework developed here could be easily adapted to incorporate derivatives, highlighting the likely benefit that could be achieved by using them in terms of reduced funding volatility. But as well as looking at the benefit of more conventional derivatives, the model developed here could be used to test the possible benefits of mortality-linked securities (see Blake *et al* (2006)), and be used to identify the appropriate design of this important and emerging area of derivatives.

7. References

Amenc, N., Martenelli, L. and Vaissie, M., (2002), Benefits and risks of alternative investment strategies, working paper, Edhec Risk and Asset Management Research Centre.

Blake, D, Cairns, A and Dowd, K, (2001), Pensionmetrics: stochastic pension plan design and value-at-risk during the accumulation phase, Insurance: Mathematics and Economics, 29, 187-215.

Blake, D. (2003), UK pension fund management after Myners: the hunt for correlation begins, Journal of Asset Management, vol. 4, pp. 32-72.

Blake, D. (2006), Pension Finance, John Wiley & Sons, Ltd, UK.

Blake, D., Cairns, A. J. G. and Dowd K., (2006) Living with mortality: longevity bonds and other mortality-linked securities, forthcoming in the British Actuarial Journal.

Driver, R. and Selvaggi, M., Understanding companies' pension deficits, Report by ABI Research Department, (March 2007).

Haberman, S., Day, D., Fogarty, D., Khorasane, M. Z., McWhirter, M., Nash, N., Ngwira, B., Wright, I.D., and Yakoubov, Y. (2003), "A Stochastic Approach to Risk Management and Decision Making in Defined Benefit Pension Schemes", British Actuarial Journal, Vol. 9 (42), pp. 493-618.

Ritchie, S. and Dumbreck, N., Letters to the Editor, Financial Times, (27th February, 2007)

The Pensions Regulator, (2006), The Purple Book: DB pensions universe risk profile.

Schneeweis, T., Karavas, V. and Georgiev, G., (2002), Alternative investments in the institutional portfolio, Working paper CISDM School of Management.

UBS Global Asset Management, Pension Fund Indicators 2006: A long-term perspective on pension fund investment, (May 2006).

Table 1**Asset classes and their respective indices**

This table contains information about the indices that are used to proxy for a range of asset class returns. Each index is used in its “total return” format and all the data were collected from Thomson Financial’s Datastream.

Asset class	Index
Cash	One month LIBOR
Gilts	Datastream UK gilt benchmark index, ten year
UK Equities	FTSE 100 index
Overseas equities	S&P 500 index (60%) Datastream Europe (ex UK) total market index (20%) Datastream Japan total market index (20%)
Sterling corporate bonds	IBOXX sterling corporate, all maturities AA index
High Yield	Lehmann Brothers, high yield index B rating
Commercial Property	UK IPD total market index
Commodities	Dow Jones AIG Commodity futures Index
Market Neutral	CSFB Tremont Market Neutral index
Global macro	CSFB Tremont Global macro index

Table 2**Asset class expected returns and correlation matrix**

This table contains information relating to the expected return and return correlations on a range asset classes. The expected return estimates were calculated using a forward-looking assessment of the likely return prospects for each asset class. The correlations between each asset class were estimated using the associated indices described in Table 1. The correlations are unconditional and were all calculated using monthly data on the indices from January 1998 to April 2007.

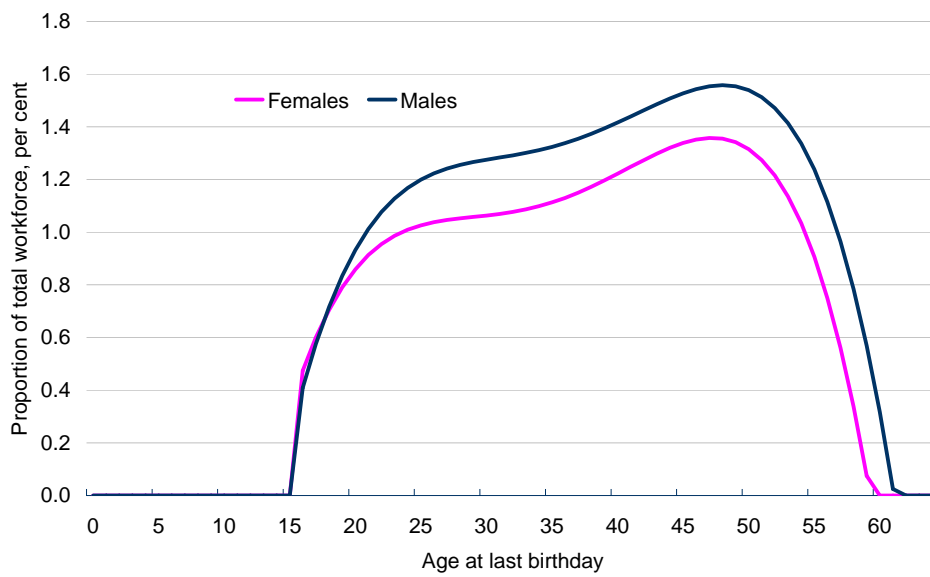
	Expected returns	Cash	Gilts	UK eq	OS eq	£ AA Corp	HY	Comm Prop	Commods	Mkt neutral	Macro
Cash	5.00%	1									
Gilts	5.25%	0.20	1								
UK Eq	8.00%	0.03	-0.20	1							
O/S Eq	8.50%	0.00	-0.20	0.84	1						
£ AA Corp	6.25%	0.14	0.91	-0.09	-0.12	1					
HY	7.50%	-0.14	-0.04	0.28	0.40	0.13	1				
Comm Prop	7.50%	-0.09	-0.05	0.18	0.19	-0.09	-0.13	1			
Commods	7.00%	-0.23	-0.04	0.05	0.16	-0.05	0.06	0.03	1		
Mkt neutral	9.00%	0.29	-0.06	0.28	0.34	-0.05	0.09	-0.03	0.10	1	
Macro	9.00%	-0.14	0.02	0.09	0.13	0.06	0.17	0.08	0.10	0.02	1

Table 3**Asset class expected returns and correlation matrix**

This table contains information regarding the asset allocation to the range of asset classes. Working from left to right the allocation to alternative asset classes – defined here as High yield bonds, commercial property, commodity futures, market neutral and macro hedge funds – gradually increases.

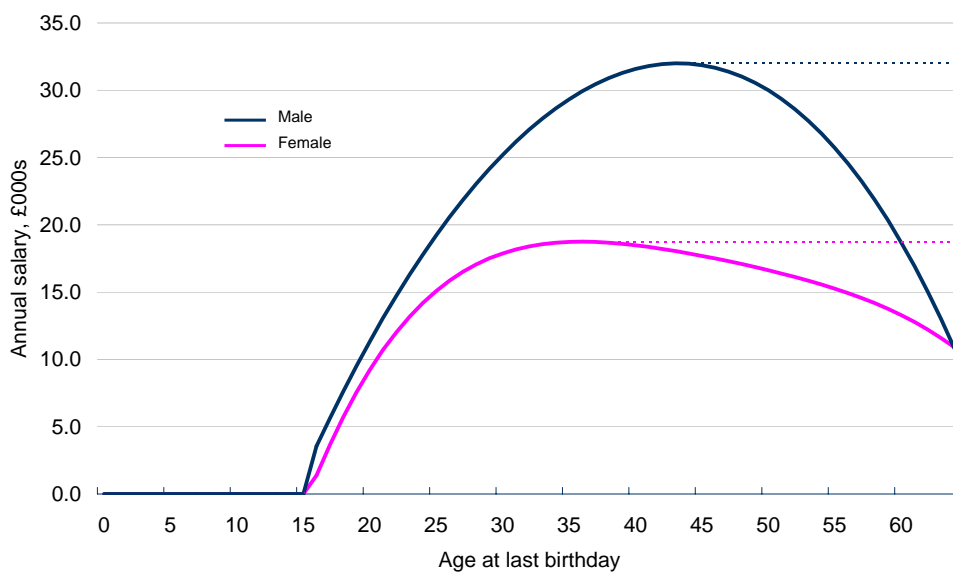
	Base	Traditional	Toe dipping	More adventurous	A bigger plunge	Dive straight in	"Yale" model
Cash	5%	0%	0%	0%	0%	0%	0%
Gilts	35%	20%	20%	20%	20%	15%	10%
£ AA Corp	0%	20%	20%	20%	20%	15%	10%
UK eq	60%	30%	25%	20%	10%	10%	10%
OS eq	0%	30%	30%	30%	30%	30%	20%
HY	0%	0%	1%	2%	4%	6%	10%
Com. Prop	0%	0%	1%	2%	4%	6%	10%
Commods	0%	0%	1%	2%	4%	6%	10%
Mkt. neutral	0%	0%	1%	2%	4%	6%	10%
Macro	0%	0%	1%	2%	4%	6%	10%
Total	100%	100%	100%	100%	100%	100%	100%

Figure 1: Age/employment profiles



Source: ONS, authors' calculations

Figure 2: Age/earnings profiles



Source: Ashe survey ONS, and authors' calculations

Figure 3

Future value of scheme liabilities of representative scheme as at 2006

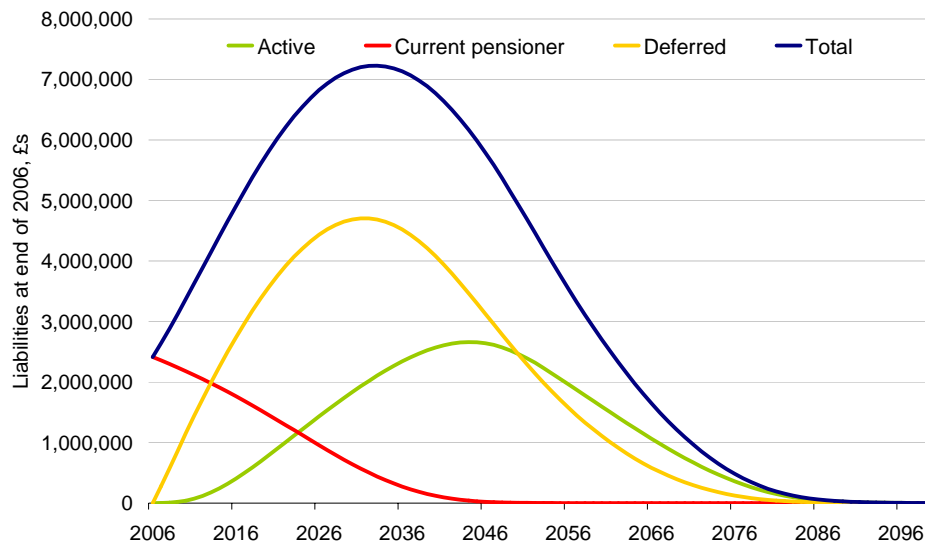




Figure 4

Present value of scheme liabilities and scheme assets through time


Base case:
expected assets, with 90% c.i.



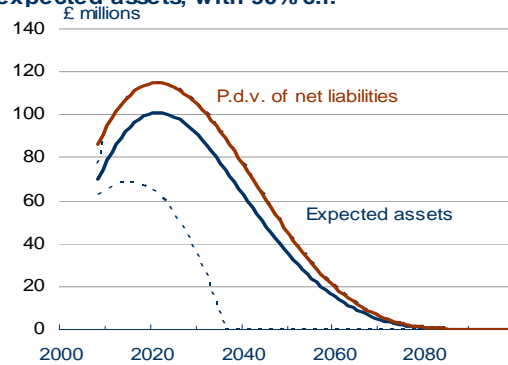
Traditional:
expected assets, with 90% c.i.




Toe dipping:
expected assets, with 90% c.i.




More adventurous:
expected assets, with 90% c.i.



A bigger plunge:
expected assets, with 90% c.i.



Dive straight in:
expected assets, with 90% c.i.



"Yale" model:
expected assets, with 90% c.i.




Figure 5
Proportion of underfunded schemes

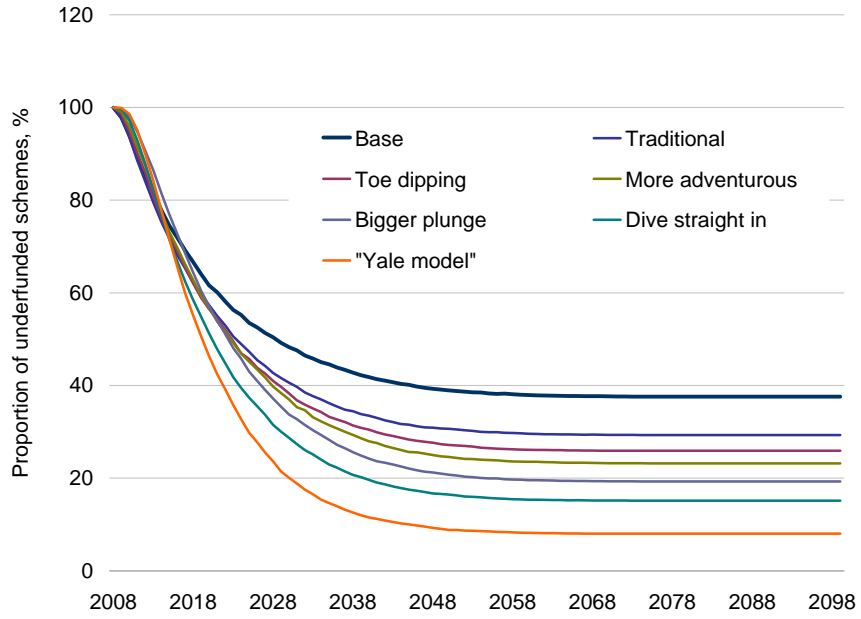


Figure 6
Average scheme funding ratios

