

Risk-Profiling Defined Benefit Pension Schemes

MICHAEL A.H. DEMPSTER, MATTEO GERMANO,
ELENA A. MEDOVA, JAMES K. MURPHY, DERMOT RYAN,
AND FRANCESCO SANDRINI

MICHAEL A.H. DEMPSTER is a professor emeritus and the founder of the Centre for Financial Research, Statistical Laboratory, University of Cambridge, and a managing director of Cambridge Systems Associates Ltd., in Cambridge, U.K.
mahd2@cam.ac.uk

MATTEO GERMANO is head of Global Investment Solutions at Pioneer Investments SA in Milan, Italy.
matteo.germano@pioneerinvestments.com

ELENA A. MEDOVA is a senior research associate at the Judge Business School, University of Cambridge, and a managing director of Cambridge Systems Associates Ltd., in Cambridge, U.K.
eam28@cam.ac.uk

JAMES K. MURPHY is a senior analyst at Cambridge Systems Associates Ltd. and a doctoral student at the Centre for Financial Research, University of Cambridge, in Cambridge, U.K.
james.murphy@cambridge-systems.com

DERMOT RYAN is an institutional portfolio manager at Pioneer Investments in Dublin, Ireland.
dermot.ryan@pioneerinvestments.com

FRANCESCO SANDRINI is head of Institutional Portfolio Managers at Pioneer Investments in Munich, Germany.
francesco.sandrini@pioneerinvestments.com

Recent market turbulence, falling valuations of equity and property, and lower long-term bond yields have highlighted again the problems of pension fund deficits. Since the start of 2008 some of the worst-affected U.K. pension schemes have lost considerable amounts of capital and the substantial falls in interest rates have lowered discount factors to result in increased liabilities across the board. The total funding situation of U.K. defined benefit (DB) plans swung from a surplus of £130.4 billion in June 2007 to a £194.5 billion deficit in December 2008 (ONS [2008b]). Among the worst-affected companies have been BT, Northern Foods, and AGA Foodservice.

The international situation is varied, owing to different regulatory environments and historical backgrounds in different countries. Countries that have had mandatory pension funds for many years exhibit the largest pension fund totals with respect to the size of their economies (Clark, Munnel, and Orzag [2006]). This is most evident in the Netherlands and Sweden where the ratios of assets to liabilities (coverage ratios) are highest. The U.K., U.S., Ireland, Germany, and Austria are lagging in terms of coverage, and a third group, including France, Italy, Turkey, China, and India, has very low coverage ratios. Accounting standards also affect the picture; standards such as FTK in Holland, FRS 17 in the U.K., and the international IAS 19 all require liabilities

to be discounted on a corporate bond yield equivalent, but they differ in terms of practice, and changing the discount rate to a risk-free rate, which would increase liabilities even more, has been mooted in a number of regions.

In considering the risks pension funds face we must look at the fund in its entirety. On the liability side, funds are exposed to interest rates, inflation, and mortality (Babel, Bornsdorf, and Kahlenberg [2008]). On the asset side, they face exposure to risk in different markets, such as fixed income, equity, and credit (Clark, Munnel, and Orzag [2006]). Hence, they operate in a multifactor environment, exposed to a wide range of asset classes and liabilities, both of which also have varying degrees of correlation among themselves. Because of this, the worst situation for pension funds is falling asset prices and lower bond yields because low yields lead to less discounting of future liabilities.

Approaches to providing solutions to the pension fund problem have developed over the last three to five years (see, for example, Deutsche Bank [2005]). In so-called first-generation products, pension funds attempted to hedge the interest rate and inflation components of their exposure using a mixture of bonds, interest rate swaps, and inflation swaps. Funds invest in a variety of solutions, mostly individually on a bespoke basis (see, for example, Goldman Sachs [2008]), but also for some smaller schemes on a pooled-fund basis, both of which

involve cash flow matching and multiyear buckets with a single instrument in each bucket.

Although this has had the effect of hedging unwanted interest rate and inflation volatility, it also hedged “good” volatility (see Dempster, Evstigneev, and Schenk-Hoppé [2007, 2008]) by locking in the funding ratio. The main consequence of single-factor liability hedging, therefore, is that it crystallizes any existing deficit (Waring and Siegel [2007]) and immunizes the scheme from good, as well as bad, volatility. This may be acceptable in schemes for which the funding ratio is already very high, but it can be a problem for underfunded schemes. The most serious consideration for any liability-hedging solution is the level of current long-term yields. When real bond yields are at a historic low, locking-in at this point in the yield cycle may be particularly undesirable.

Recent approaches to the problem are more holistic and attempt to manage liability exposure while seeking to add value within a given risk budget. Second-generation solutions tend to focus on maximizing the excess return above a liability benchmark at some specified level of risk (Babel, Bornsdorf, and Kahlenberg [2008]).

In this article, we will present a new type of solution for defined benefit pension schemes. The approach we adopt uses Monte Carlo generation of asset returns, asset prices, and liabilities of a defined benefit pension scheme over a wide range of economic conditions. We employ dynamic stochastic optimization to determine the optimal asset allocation and employer contribution rates that will enable the scheme to achieve a desired funding ratio within a given time horizon while respecting the trustees’ risk appetite.

Our approach is *dynamic* and is based on an active asset allocation strategy with periodic portfolio rebalancing. We view this as superior to a sequential *static* Markowitz surplus-type strategy for several reasons, such as the dynamic nature of the liabilities and the fact that dynamic asset allocation can cope with the flow of capital into and out of the scheme in the form of contributions and benefit payments. Also, static asset allocation relies on Gaussian return distributions for assets, whereas dynamic asset allocation based on Monte Carlo is free to adopt any appropriate return distributions.

The article is organized as follows. The next section briefly discusses the concepts behind the stochastic optimization techniques employed in the article. An overview of the asset return and liability generation using CASM, Pioneer Investments’ proprietary corporate simulator, is summarized in the third and fourth sections.


In the fifth section the multi-objective dynamic stochastic optimization problem of restoring a DB pension scheme to a target coverage ratio is formulated and solved. The sixth section gives the results of two simplified example schemes—one underfunded and the other well funded. The last section concludes.

OVERVIEW OF DYNAMIC STOCHASTIC PROGRAMMING

Our approach to pension fund management uses dynamic stochastic programming to select allocations that are optimal with respect to fund liabilities and suitable measures of underfunding risk (Dempster, Germano, Medova, and Villaverde [2003] and Dempster, Germano, Medova, Rietbergen, Sandrini, and Scrowston [2006, 2007]). *Dynamic stochastic programming* (DSP) involves simulating economic factors, asset returns, and liabilities forward over a number of scenarios. Using simulated market scenarios allows any distribution to be used for each of the asset returns or liabilities according to the model for the dynamics of these quantities.

The scenarios are arranged in a tree structure (shown in Exhibit 1) and decisions are made at points where the tree branches. Each of these decisions is optimal with respect to all the simulated evolutions of the asset returns and liabilities that could occur after the decision point. As the *root node* corresponds to the present moment, the decision made there is the one that is to be implemented. This means that

EXHIBIT 1
Scenario Tree with Single Scenario



Note: Single scenario is the dashed line.

the root-node decision is robust with respect to all generated future scenarios representing future market conditions.

In DSP, scenarios are used as a way of representing current uncertainty about the future. Since only the decision at the root node is to be implemented, other decisions further in the (simulated) future can be viewed simply as a what-if model of rational behavior at those future points in time. In real applications, future decisions are made by re-running the simulation and optimization in the future after incorporating the additional information that has become available through the passage of time. In the DSP Monte Carlo approach used in this article, all scenarios are equally weighted, meaning that the probability of any particular one is inversely proportional to the total number of scenarios at the particular point in time being considered. To enhance the computational efficiency of solution it is possible to generate the tree iteratively based on the expected value of perfect information at each of the decision points (Dempster [2006]), increasing the branching of the tree to the planning horizon specifically in areas that have more influence on the objective function.

Clearly, a crucial step in the DSP approach is the forward simulation of the quantities of interest—in this case, economic factors, asset returns, and liabilities. Thus, a requirement of the scenario generator in the present context is that it is possible to generate scenarios corresponding to a wide range of market conditions for both asset returns and liabilities. In the subsequent sections we will describe the simulation of the quantities used in our model.

ASSET RETURN SIMULATION

Asset simulation is based on Pioneer Investments' proprietary corporate simulator CASM (Cascade Asset Simulation Model). Cascade simulators for asset-price-return simulation were introduced by Mulvey [1996]. Another implementation for Pioneer Investments is given in Dempster et al. [2003] and Arbileche-Gréla and Dempster [2005]. In this implementation, the complex nonlinear relationship between parameters of the major currency region models is approximated by analytical expressions describing the drift parameters, and the currency regions are linked by exchange rates.

CASM is a modification and extension of these simulators reflecting the overall view of Pioneer Investments and linking different risk factors across divisions of the firm.

CASM

CASM is based on a hierarchical econometric model with European Union GDP growth and inflation considered as the basic driving forces from which further economic factors, such as interest rate term structure and asset prices, are derived, as shown in Exhibit 2. The economic and financial risk factors are modeled with parameters and correlations fitted to quarterly data in a three-level *cascade* fashion, as follows:

Level 1: Macro factors: GDP cycle, GDP trend, oil cycle, and real short rates.

Level 2: Cyclical financial variables and term structure.

Level 3: Equities, credit spreads, and FX.

The pension scheme asset classes used in the examples in the sixth section of this article is a selection proxied by nominal and inflation-linked bond indices of various durations, a money market fund index, and an equity index. For this set of asset classes, the key quantities that determine their returns are interest rates and term structure, inflation, and equity performance.

The asset classes in Exhibit 3 were simulated for potential investment using CASM. The complex nature of the nonlinear interaction between asset returns and economic variables in CASM result in dynamic interdependencies between the variables. For each asset class a proxy index of similar instruments

EXHIBIT 2

Cascade Asset Simulation Model (CASM)




EXHIBIT 3

Asset Classes and Related Indices

Asset class	Code	Index
Money Market	EUR300M	3 month EURIBOR rate
1–3 year duration Euro bonds	EUM1TR	Bloomberg 1–3 year Euro bond market tracker
3–5 year Euro bonds	EUM2TR	Bloomberg 3–5 year Euro bond market tracker
5–7 year Euro bonds	EUM3TR	Bloomberg 5–7 year Euro bond market tracker
7–10 year Euro bonds	EUM4TR	Bloomberg 7–10 year Euro bond market tracker
10+ year Euro bonds	EUM5TR	Bloomberg 10+ year Euro bond market tracker
1–3 year inflation indexed bonds	BEIG1T	Barclays Euro Govt Inflation linked bond tracker for all maturities
5 year inflation indexed bonds	BEIG2T	Barclays Euro Govt Over 5 Year
5–7 year inflation indexed bonds	BEIG3T	Barclays Euro Govt 1–10 Year
10+ year inflation indexed bonds	BEIG4T	Barclays Euro Govt Over 10 Year
15+ year inflation indexed bonds	BEIG5T	Barclays Euro Govt Over 15 Year
European equity	DJST	Dow Jones EuroStoxx 50 stock index

(see Exhibit 3) was used to assess the quality of the simulated scenarios.

To illustrate the complex nature of the output from CASM, forward simulations of the corresponding asset class returns have been performed quarterly from Q1 2006 to Q4 2014. Representative scenarios are shown in Exhibit 4 in the form of fan charts. These charts illustrate the distribution of the scenarios across time as a fan plot, with the intensity of the shading declining through quantiles either side of the median (so that, for example, the 20th and 80th percentile values across all scenarios are shaded the same). The median scenario, interquartile range, and envelope (the faint dotted line at the edge of the fans that shows the most extreme value in any scenario in the given period) are marked for clarity.

These diagrams also allow us to make some assessment of the plausibility of our asset simulations with respect to historical data up to August 2008. In almost all cases the envelope of the simulated scenarios includes realized returns over the simulation period seen to that date (the exceptions are in short-duration nominal bonds where recent large movements have gone beyond the scenario envelope). The median of the simulated scenarios also fits with the observed returns during this period. In all cases the volatility of the scenarios is on the same scale as the historical return volatility.

MODELING LIABILITIES

Pension fund liabilities arise both from the number of members in the scheme and from the cost of purchasing annuities now and in the future to fund members' retirement pensions. The sizes of the annuities that must be purchased depend on the final salary of the scheme members at retirement and, hence, on interest and inflation rates. The simplified model of the total defined benefit liabilities we use for illustration in this article is based on Cairns [2004]. The techniques required to consider mortality risk may, however, be incorporated in the DSP framework (see, for example, Medova, Murphy, Owen, and Rehman [2008]).

Scheme Structure Model

The number of members of different ages in a scheme, M^t , is modeled starting with an *initial age profile* for the scheme,


$$M^0 = (M_{x_0}^0, M_{x_1}^0, \dots, M_{x_n}^0) \quad (1)$$

where M_x^t is the number of members in the scheme at time t of age x .

This evolves deterministically over time with members moving from one age band to the next band each year,

EXHIBIT 4

Fan Charts of Simulated Quarterly Asset Returns for Investment Assets




$$M_{x_i}^{t+1} = M_{x_{i-1}}^t \quad i = 1, \dots, n \quad (2)$$

All members are assumed to retire at age x_n . Our simplified model only allows members to join the scheme at age x_0 and does not account for death in service or early retirement. It would, however, be

straightforward to add these features to our model and to allow members to join at any age. The process by which members join the scheme, $\mathbf{M}_{x_0}^t$, can be specified by a fund's own model or modeled using an appropriate stochastic process.¹

EXHIBIT 4 (Continued)



Salary Model

A model of the salaries of scheme members is important for two reasons: 1) contributions from scheme members are made as a fixed proportion of their salary, and 2) being a *final salary* pension scheme, liabilities depend on scheme members' final salaries. The first consideration means that it is necessary to have a complete salary model rather than just a model of salary at retirement.

We use a simple model for salary evolution. Real salaries increase with age each year, with the proportional increase fixed for each particular age. Salaries also increase with inflation, so that their real value is not eroded over time.

We start with an *initial profile of salaries*,

$$S^0 = (S_{x_0}^0, S_{x_1}^0, \dots, S_{x_{n-1}}^0) \quad (3)$$

We use this to define the (*real*) *age-related salary increase* from age x_i to x_{i+1} that will be used throughout the simulation as $(S_{x_{i+1}}^0/S_{x_i}^0)$.

Nominal salary evolution is, therefore,

$$\mathbf{S}_{x_i}^{t+1} = \left(\frac{S_{x_i}^0}{S_{x_{i-1}}^0} \right) \mathbf{c}(t, t+1) \mathbf{S}_{x_{i-1}}^t \quad (4)$$

where $\mathbf{c}(t, t+1)$ is the CPI inflation factor from time t to $t+1$. Note that since inflation is a simulated stochastic quantity the resulting salaries in Equation (4) are scenario dependent.

The expected *final salary* of a scheme member of age x_i can be written as

$$\mathbf{F}_{x_i}^t = \mathbb{E}_t(\mathbf{c}(t, t+n-i)) \mathbf{S}_{x_i}^t \left(\frac{S_{x_n}^0}{S_{x_i}^0} \right) \quad (5)$$

Because the inflation process from t to $t+n-i$ is not known at t , we must take its conditional expectation (over forward scenarios) in Equation (5) in order to estimate the final salaries of scheme members at time t , which is necessary for calculating liabilities.

As with the scheme structure model, it would be completely straightforward to replace this salary model with a more nuanced model of salary evolution corresponding to a fund's actuarial forecasts.

Contributions

Contributions are made to the scheme from two sources—employees and the employer to whom the scheme belongs. Employee contributions are modeled as a fixed proportion of salary.² Employer contributions are a variable proportion of the total salary roll, typically not greater than 20%, and the exact level is determined at the optimization stage.

Annuity Model

Pension benefits to retirees are generally funded through the purchase of *life annuities*, which are life-linked products that pay a constant (or inflation-linked) payment at some regular interval until the death of the annuity holder. Many schemes have retirement payments that are linked to consumer price inflation (CPI) over the lifetime of the annuity.

We use a simplified model of mortality, assuming that all scheme members live for a fixed number of years from retirement. This allows us to price the annuities that fund their pensions as the current present value of all future payments for this deterministic period (Milevsky [2006]). Though such deterministic mortality might appear to be a major simplification, in the absence of any currently traded liquid asset that is capable of hedging such mortality risk, this simplification will not have a qualitative effect on scheme asset allocations; allocations will concentrate on hedging inflation and the duration exposure of liabilities. It is a consequence of Jensen's inequality (Milevsky [2006, p. 116]) that the price of an annuity calculated with stochastic mortality is strictly less than the price of an annuity with a fixed length equal to the expected life remaining. Our estimates of liabilities based on fixed-length annuities will therefore be pessimistic when compared to those made with a stochastic mortality model.

Thus, to price the life annuities used to fund retiree benefits, we only need to price an annuity with constant *tenor* (i.e., payment length), where the tenor is the expectation of employee survival after retirement. The time t price of a *non-inflation-linked unit annuity due* (i.e., one paying a single unit of currency each year with immediate initial payment) is

$$\ddot{\mathbf{a}}_{\overline{n}|}(t) = \sum_{i=0}^{n-1} \mathbf{v}(t, t+i) \quad (6)$$

where

n is the tenor of the annuity in years, and $\mathbf{v}(t, t+i)$ is the time t value of a unit cash flow occurring at time $t+i$.

The rates $\mathbf{v}(t, \tau)$ depend on the simulated term structure of interest rates given by

$$\mathbf{v}(t, \tau) = (1 + \mathbf{r}(t, \tau))^{-(\tau-t)} \quad (7)$$

where $\mathbf{r}(t, \tau)$ is the interest rate on bonds maturing at $\tau > t$ taken from the yield curve simulated at t . These rates are scenario dependent because the term structure of interest rates is simulated on each scenario, thus making the annuity prices in Equation (6) scenario dependent. If the annuity is inflation linked then the formula for a *real* unit annuity becomes

$$\ddot{a}_{\overline{n}|}^{\text{real}}(t) = \sum_{i=0}^{n-1} v_{\text{real}}(t, t+i) \quad (8)$$

where $v_{\text{real}}(t, t+i)$ is the *real* time t value of a unit cash flow occurring at time $t+i$.

The rate v_{real} is defined as

$$v_{\text{real}}(t, T) = (1 + r_{\text{real}}(t, T))^{-(T-t)} \quad (9)$$

where $r_{\text{real}}(t, T)$ is the real (as opposed to nominal) interest rate at maturity T taken from the simulated real yield curve at time t .

Total Liabilities

In the U.K., FRS17 demands that the total liability of a pension scheme be calculated as the present value of all future expected payments, taking into account only payments from service given to date, but calculating the benefits on the basis of the inflation-adjusted expected final salary.

According to this rule (and on the basis of our simplified scheme and salary structures) the member liability to the scheme at time t of a member of age x_i , who has provided i years of service and has another $(n-i)$ years until retirement, is

$$\underbrace{\frac{(x_i - x_0)}{60} \mathbf{F}_{x_i}^t}_{\text{expected liability based on years of service}} \cdot \underbrace{v(t, t+n-i) \ddot{a}_{\overline{n}|}^{\text{real}}(t+n-i)}_{\text{current value of future annuity}} \quad (10)$$

The total scheme liability at time t is the sum of liabilities for all scheme members

$$\mathbf{L}_t = \sum_{i=0}^{n-1} M_{x_i}^t \frac{(x_i - x_0)}{60} \mathbf{F}_{x_i}^t \cdot v(t, t+n-i) \ddot{a}_{\overline{n}|}^{\text{real}}(t+n-i) \quad (11)$$

The fund liability is scenario dependent and is shown for nominal annuities, but these could be substituted by real annuities using Equation (8) if required. Using a deterministic model of fund membership and salary structure means that the stochastic component of fund liabilities is driven by inflation and interest rates; adding a stochastic component to membership or salary structure would add a further random element to fund liability. Depending on

the model used, this random component might be difficult to hedge using the assets currently available in the market.

OPTIMIZATION PROBLEM

Annuity prices, and therefore fund liabilities, are heavily dependent on interest rates and, when the annuities are inflation linked, on inflation rates. Since these basic processes underlie our asset return models, the fund's liabilities are linked to its returns. The evolution of the fund's portfolio must be constructed in such a way as to optimize its exposure to various instruments in order to match its liabilities. This is the basis of the asset-liability management problem in the absence of mortality risk.

To solve this problem optimally, we must trade off considerations of three conflicting objectives in each time period—maximizing fund wealth, minimizing shortfall relative to a target wealth, and minimizing employer contributions. We construct a target wealth for the fund for each period in terms of a *target funding ratio*, which is a desired ratio of current assets to liabilities at the end of the planning period. The target funding ratio in each period, TR_t , is set to vary linearly between the initial funding ratio, TR_1 , and the final target funding ratio, TR_T , at the amortization horizon, T ; that is,

$$TR_t = \frac{T-t}{T-1} TR_1 + \frac{t-1}{T-1} TR_T \quad t = 2, \dots, T-1 \quad (12)$$

The target wealth of the fund is then simply the target funding ratio multiplied by the total liability of the fund at that period,

$$\mathbf{TG}_t = TR_t \mathbf{L}_t \quad t = 1, \dots, T \quad (13)$$

Note that the fund liabilities, $\mathbf{L}_t$, are stochastic, so that, unlike the target funding ratio, the target wealth is scenario dependent.

Fund Shortfall and Surplus

The shortfall of the fund $\mathbf{w}_t^-$ is defined as the amount by which the fund wealth $\mathbf{w}_t$ is below the target wealth, $\mathbf{TG}_t$, or zero if there is a surplus. The fund surplus $\mathbf{w}_t^+$ is defined by the amount by which the fund wealth is above the target wealth or zero if there is a shortfall. Clearly these possibilities are mutually exclusive on each scenario

at any given point of time (i.e., the fund is either below or above its current target). The relations between the variables are therefore

$$\begin{aligned} \mathbf{w}_t^+ - \mathbf{w}_t^- &= \mathbf{w}_t - \mathbf{T}\mathbf{G}_t \\ \mathbf{w}_t^+ &\geq 0 \\ \mathbf{w}_t^- &\geq 0 \quad t = 1, \dots, T \end{aligned} \quad (14)$$

Objective

Pension funds would ideally like to reach their target funding ratio while minimizing both their exposure to risk and to the level of additional contributions made by the employer to whom the scheme belongs. The objective function to be maximized reflects these conflicting objectives as

$$\mathbb{E} \left(\sum_{t=1}^T \alpha \mathbf{w}_t^+ - \beta \mathbf{w}_t^- - \gamma (1-d)^{t-1} \mathbf{c}_t \mathbf{S}_t \right) \quad (15)$$

where

$\mathbf{c}_t$ is the proportion of the total salary roll contributed by the employer at time t ;
 d is the discount rate per period used by the employer;
 $\mathbf{S}_t$ is the total salary roll at time t ;
 T is the number of periods over which the dynamic optimization takes place; and
 α , β , and γ are weights on the objective terms, which set the balance of importance between the competing objectives.

The first term in the objective aims to maximize total expected surplus over periods when the fund exceeds its target, the second to minimize total expected shortfall over periods when the fund falls short of its target, and the third to minimize total expected employer contributions over all periods.

All these quantities are in units of currency. Shortfall and surplus would be expected to be comparable in magnitude and might be in the region of 10% of the fund wealth. We can roughly determine the size of the contributions in terms of fund wealth by assuming full funding and then equating fund wealth with fund

liability defined in terms of salary. For full funding the liability is approximately equal to the value of, say, a 15-year annuity of two-thirds of the final salary for all retired scheme members, and for a fund in the steady state is something around two-thirds of the total salary roll of employed scheme members. The result is a fund wealth in the region of 10 times the total salary roll, and because contributions are in the region of 0%–20% of this figure (ONS [2009a]), this means that the contribution term will be in the region of 1% of the fund wealth. So, the objective-term values have a magnitude ratio of roughly 10:10:1 for surplus, shortfall, and contributions, respectively.

If we assume that varying decisions have a proportionally similar effect on the objective terms, this offers some indication as to the appropriate relative magnitudes of the objective-term weights. This is a significant assumption, but it can at least serve as a heuristic to help us in the task of finding appropriate parameter values within the parameter space of the model. Because fund shortfall is much more important than surplus, which is not beneficial per se, β should be much greater than α . The contribution term is of similar importance to the shortfall, but can be expected to be around 10 times smaller than the shortfall term, so γ should be roughly 10 times the size of β ,

$$\begin{aligned} \alpha &\ll \beta \\ \gamma &\approx 10\beta \end{aligned} \quad (16)$$


We also discount the employer contributions by some amount to reflect the fact that future payments to the scheme are less costly (and thus more desirable) for the employer than current ones. The discount rate d used by the employer depends on the company's cost of capital, so it would be expected to be in the region of 10% a year.

Cash Flow Constraints

The cash flow model represents the basic movement of value in the fund and ensures proper accounting of cash flows into, out of, and within the fund. It is represented by Exhibit 5, in which circles represent stores of value and arrows represent paths along which value can be transferred.

EXHIBIT 5

Cash Flow Diagram for Defined Benefit Pension Fund



This leads to the *cash flow constraint*,

$$z_t = (1 + r_t^{\text{cash}})z_{t-1} + \mu_E S_t + c_t S_t + \sum_{a \in A} x_{a,t}^- - \sum_{a \in A} x_{a,t}^+ - B_t \quad a \in A, t = 2, \dots, T \quad (17)$$

where

A is the set of assets from which the portfolio can be composed,

B_t is the payment of employee benefits (annuity purchases) at time t ,

r_t^{cash} is the interest rate on bank cash at time t ,

$x_{a,t}^-$ is the *sale* by value of asset a at time t ,

$x_{a,t}^+$ is the *purchase* by value of asset a at time t ,

z_t is the holding of bank cash at time t , and

μ_E is the (fixed) proportion of salary contributed by employees.

This also leads to the *asset balance constraints*, which account for the asset holdings and their accrual of value due to returns,

$$x_{a,t} = x_{a,t-1}(1 + r_{a,t}) + x_{a,t}^+ - x_{a,t}^- \quad a \in A, t = 2, \dots, T \quad (18)$$

where

$x_{a,t}$ is the fund's holding of asset a at time t , and $r_{a,t}$ is the return on asset a from time $t-1$ to t .

These two constraints do not apply in the first period (since z_0 and the $x_{a,0}$ are not defined) and are replaced in that period by the initial cash flow constraint,

$$z_1 = W_0 + \mu_E S_1 + c_1 S_1 + \sum_{a \in A} x_{a,1}^- - \sum_{a \in A} x_{a,1}^+ - B_1 \quad (19)$$

where W_0 is the initial wealth of the fund, and by the initial asset balance constraints,

$$x_{a,1} = x_{a,1}^+ - x_{a,1}^- \quad a \in A \quad (20)$$

which simply state that the initial asset holding is equal to the amount of an asset bought less the amount of an asset sold at Time 1.

For simplicity, we abstract from fund management fees, transactions costs, and turnover constraints, but these may easily be included (see, for example, Dempster et al. [2006, 2007]).

Total Wealth

The total wealth of the fund is defined as the sum of all asset holdings; that is,

$$\mathbf{w}_t = \sum_{a \in A} \mathbf{x}_{t,a} \quad t = 1, \dots, T \quad (21)$$

Employer Contribution Rate

The level of employer contributions is constrained to lie between upper and lower limits ($c_{\min}$ and $c_{\max}$, respectively) specified by the user; that is,

$$c_{\min} \leq \mathbf{c}_t \leq c_{\max} \quad t = 1, \dots, T \quad (22)$$

Short-Selling Constraint

In this model, we restrict the fund to long-only positions, which means that all asset holdings must be non-negative,

$$\mathbf{x}_{a,t} \geq 0 \quad a \in A, t = 2, \dots, T \quad (23)$$

Cash Constraint

Cash is constrained to be non-negative, so that the fund is not allowed to borrow from the bank at the cash rate,

$$\mathbf{z}_t \geq 0 \quad t = 1, \dots, T \quad (24)$$


Solution

In order to make a large-scale dynamic stochastic programming problem such as ours computationally tractable, we must limit the number of scenarios and branch points at which portfolio rebalances are undertaken. The CASM simulator produces asset returns at quarterly time steps, however, and this is used as the basic frequency of accounting and decisions in the model. In order to limit the number of scenarios while still using quarterly simulation data and allowing quarterly portfolio and contribution adjustments, we adopt a tree that only branches at a limited subset of stages, as shown in Exhibit 6.

The difficulty with this two-time-scale scenario tree structure without further enhancement is that at periods between major rebalance points the exact evolution of the scenario up to the next major rebalance point is visible to the optimizer. In order to prevent the optimizer exploiting this information, we insist that it makes these portfolio adjustment and contribution assignment decisions for the nonbranching periods at the preceding

EXHIBIT 6

2.2.2 Tree with Ten Stages and Three Branching Stages



Note: 2.2.2 tree is shown in black.

branching stage (major rebalance point), where there is still uncertainty as to which scenario will be realized.

This introduces the further problem that, at the preceding branch point, the exact value of an asset at forthcoming stages is not known. Thus the exact value of purchases or sales of the asset are unknown, which could result in a surfeit or shortage of bank cash on a particular scenario. In fact, it is possible to exclude one of these mutually exclusive possibilities by restricting the sign of the cash variable, and so we chose to only allow positive cash holdings, enforced by the non-negativity constraint in expression (24). Overcoming this effect is the reason for introducing noninvestable bank cash, which acts as a slack variable able to absorb excess value between major rebalance points. Bank cash holdings tend to grow slowly between major rebalance points and return to zero at the major rebalance points. Bank cash is not intended to be an investment asset, so we make the cash interest rate uncompetitive with the returns on other assets (we have used a zero cash interest rate in our examples).

We formulate the model and solve the example problems using tools from the *STOCHASTICS*TM suite for dynamic stochastic programming problem formulation, analysis, and visualization (CSA [2008]).

EXAMPLES

We present two sample problems. In the first, we have a severely underfunded scheme that aims to become

fully funded within 10 years. In the second, we have an almost fully funded scheme that aspires to improve its funding ratio over 15 years. We present two solutions for each scheme, one with variable employer contributions and another with employer contributions fixed to a constant proportion of the total salary roll.

Underfunded Scheme

We consider a hypothetical underfunded pension scheme with a joining age of $x_0 = 25$ and retirement age of $x_n = 65$. The initial age profile is

$$M^0 = (20, 20, \dots, 20) \quad (25)$$

and the joining rate is taken to be constant, $M_{x_0}^t = 20$, throughout the problem. The initial salary profile is

$$S^0 = (20, 21, \dots, 59) \quad (26)$$

where the salaries are given in thousands of euros. The annuities provided for retiring scheme members are taken to have a *constant* tenor of 15 years with inflation-linked payments.

The initial liability in Q2 2007 is calculated to be €150,235,000 and the initial fund wealth is €125,000,000 from which a payment of €10,439,000 is immediately made to pay for annuities for retirees at Time 0. Thus the initial wealth is €114,561,000, giving the scheme an initial funding ratio of 76.37%.

In the first solution, employer contributions are constrained to lie between 0% and 20% of the total payroll. This range covers the realistic range of pension contributions made by U.K. employers; in 2007, 40.7% of U.K. schemes had employer contributions between 8% and 15%, while 58.7% had contributions above 15% (ONS [2009a]). The mean was 15.6%.

In the second solution, contributions are fixed to a constant 12%. This would place the scheme in the 24th percentile of U.K. defined benefit schemes in terms of employer contributions in 2007 (ONS [2009a]).

Employee contributions are fixed at 5% of salary, which places the scheme around the 21st percentile of U.K. defined benefit schemes in terms of the 2007 employee contribution rate (ONS [2009a]), although just above the mean of 4.9%. As with employer contributions, the distribution of employee contributions is negatively skewed.

Model Parameters

The asset return and liability simulation is quarterly with a 10-year horizon from Q2 2007, quarterly portfolio rebalancing, and major rebalancing at the branching times. We use a (50.8.4.4) regular branching tree with branching times of 0, 2½, 5, and 7½ years and the investable asset classes are those described in the third section. The aim is to achieve a fully funded scheme within 10 years, giving a target funding ratio of 100%.

The optimization uses the parameters $\alpha := 10^{-6}$, $\beta := 0.1$, and $\gamma := 1$, which fit closely with the values in expression (16) determined heuristically in the preceding section. By using such a small value of α relative to β we are saying that the negative impact of €1 of shortfall is equivalent in magnitude to the positive impact of €100,000 of surplus. In other words, we are effectively requiring the optimizer never to build up surplus above the target on some scenarios at the expense of increasing shortfall on other scenarios. We use a discount rate for employer contributions of 6.5% a year.

Results

In Exhibit 7, we chart the average fund wealth across all scenarios along with its target and liability. The drop in wealth at every fourth step is due to the annual purchase of annuities for retiring members. Initially, the expected wealth is far less than the liability, but by the amortization horizon the expected wealth has reached the target in both cases.

In Exhibit 8, we chart the average level of employer contributions. For the problem with variable contributions, initial employer contribution rates are at their maximum permitted level. As the fund starts to surpass its funding target, the contributions drop steadily from 2½ years into the problem. The pattern of contributions seen in the variable-contribution optimizations makes intuitive sense, because the optimizer chooses to take employer contributions early when they can contribute to fund growth.

In Exhibit 9, we show the (expected) dynamic asset allocations averaged across scenarios. In both instances, the initial asset allocations are largely made up of long-maturity inflation-linked bonds. This is because all scenarios start with a small shortfall to the target that the optimizer must reduce but is penalized heavily for increasing. As well-performing scenarios rise above the target they can

EXHIBIT 7

Expected Wealth, Liability, and Funding Targets with Variable and Constant Employer Contributions





EXHIBIT 8

Expected Employer Contribution Rate with Variable and Constant Employer Contributions



afford to increase their exposure to equity, which can provide them with enhanced returns. The bulk of the assets are still held in inflation-linked bonds, which makes sense in light of the fund's inflation-linked liabilities. Notice that the final allocation contains a high proportion of long-maturity inflation-indexed bonds that reflects the long maturity of many of the promised payments and an attempt to hedge exposure to future inflation. The allocation for both problems is quite similar, although the constant-contribution problem contains slightly more equity in the first 2½ years to compensate for lower contributions during this period.

Exhibit 10 shows the evolution of the fund's shortfall (against the target) across all scenarios. As with the fan plots of the asset simulations in the third section, these plots show the distribution of the fund shortfall throughout

the optimization, with the median, interquartile range, and scenario-envelope marked.

The two charts are fairly similar, with variable contributions leading to lower early shortfalls. With constant contributions, the range of the extreme scenarios is increased somewhat as the fund is unable to use variable contributions to ameliorate funding difficulties or to reduce contributions when funding levels are high.

The range of fund shortfall across the scenarios can also be influenced using the α and β parameters; increasing the emphasis on minimizing shortfall (increasing β relative to α) will decrease the range of fund wealth, whereas increasing the emphasis on wealth will increase the range of fund wealth, because more risk will be taken in order to increase expected return.

EXHIBIT 9

Expected Dynamic Asset Allocation with Variable and Constant Employer Contributions





EXHIBIT 10

Fan Plots of Fund Shortfall across Scenarios with Variable and Constant Employer Contributions



The variable- and constant-contribution cases of the underfunded problem took 149s and 137s to solve, respectively, using a dual Xeon 3 GHz computer.

Well-Funded Scheme

We now consider the case of a fund with a healthier funding ratio that seeks to achieve a high funding ratio over a longer 15-year horizon from Q2 2007. Again, we set $x_0 = 25$ and $x_n = 65$. The initial age profile is

$$M^0 = (25, 25, \dots, 25) \quad (27)$$

with a constant joining rate, $M_{x_0}^t = 25$, at all times.

The initial salary profile is

$$S^0 = (25, 26, 25, 27.5, \dots, 73.75) \quad (28)$$

in thousands of euros. The annuities purchased to fund retirees' pensions are inflation linked with a constant tenor of 15 years.

In this example, the initial liability is €234,743,000, and after an initial lump-sum annuity payment of €16,311,000 from the fund value of €235,000,000, the initial wealth is €218,689,000, giving a funding ratio of 93.16%.

In the first solution, employer contributions to the scheme are constrained to lie between 0% and 20% of the total payroll. In the second solution, they are fixed to be a constant 7%, which would place the scheme around the 4th percentile of U.K. pension schemes for employer contributions in 2007 (ONS [2009a]). Employee contributions are fixed at 5% of salary.

Model Parameters

We use a (50.8.4.2.2) regular branching tree with branching times of 0, 3, 6, 9, and 12 years. The simulation is quarterly with a 15-year horizon, quarterly portfolio rebalancing, and major rebalancing at the branching times. The aim is to achieve a well-funded scheme within

15 years, with a target funding ratio of 120%. As in the underfunded case, the optimization uses the parameter values $\alpha := 10^{-6}$, $\beta := 0.1$, and $\gamma := 1$. Again, we use an employer discount rate of 6.5% a year.

Results

Exhibit 11 shows the liability and wealth for this scheme. The fund wealth starts just below the initial liability and for both constant and variable contributions increases to a final level that is above 120% of the fund's liability at the 15-year amortization horizon. In both regimes the expected fund wealth considerably exceeds the target wealth. As will be seen in the surplus plots

EXHIBIT 11

Expected Wealth, Liability, and Funding Targets with Variable and Constant Employer Contributions





EXHIBIT 12

Expected Employer Contribution Rates with Variable and Constant Employer Contribution Rates



below, this reflects the fact that the fund is above the funding target in nearly all scenarios.

Exhibit 12 charts the average employer contribution rates. In this case, we see a similar pattern to the underfunded case, with the variable-contribution fund using early contributions to reach the target funding ratio, and then reducing them when a high level of funding has been achieved. Total employer contributions made under the variable-contribution regime are €40,900,000 (present value = €34,800,000), lower than the €61,000,000 (present value = €37,400,000) made in the fixed-contribution rate scheme, yet the expected fund wealth is just €3,100,000 (present value = €940,000) higher at the

problem horizon in the constant-contribution scheme. In reality, the value of early contributions to a pension scheme is determined by the rate of discounting used by the employer. If the company uses a high discount rate (indicating, for example, a high cost of capital) early contributions will appear less attractive.

Exhibit 13 presents the dynamic asset allocation averaged across scenarios. The asset allocations are similar to those seen in the underfunded case, with the fund relying on inflation-linked bonds to make up the bulk of its assets, in line with its liabilities. As funding levels increase, the fund can afford to allocate more of its wealth to equity. The allocations for the two contribution

EXHIBIT 13

Expected Dynamic Asset Allocation with Variable and Constant Employer Contributions





EXHIBIT 14

Fan Plots of Fund Shortfall Across Scenarios with Variable and Constant Employer Contributions



schemes are similar, with the constant-contribution fund taking on more risk through a higher equity holding in the first three years.

Finally, the fan plots of the fund shortfall in Exhibit 14 show that under both regimes the fund is expected to stay above its funding target in almost all scenarios. As with the underfunded case, the constant-contribution regime initially falls below the target in many scenarios because variable contributions cannot be used to make up the shortfall.

The variable- and constant-contribution cases of the well-funded problem took 362s and 333s to solve, respectively, using a dual Xeon 3 GHz computer.

CONCLUSION

In this article, we have described a strategic-level dynamic stochastic optimization model that can be used to advise underfunded defined benefit pension schemes on best practice for returning the fund to solvency and long-term stability. We have presented an overview of the dynamic stochastic programming techniques involved and briefly described the nature of Pioneer Investments' proprietary CASM simulator from which the asset class returns and pension scheme liabilities were generated. The stochastic optimization model was set out precisely and its solution by linear programming discussed. To illustrate the approach, two examples of defined benefit schemes using simple, conservative fund-liability models were presented and discussed. The optimal dynamic asset allocations of these examples reflect the motivation of second-generation liability-driven investment schemes discussed in the introduction. Although the final salary scheme models used in our examples are simple, more complex models can be incorporated with little extra effort into the system described. Most actuarial assessments used in practice (see, for example, Clark, Munnel, and Orzag [2006]) can be modeled for this purpose.

ENDNOTES

¹We use bold-face type throughout the article to denote random scenario-dependent entities.

²With a mean of 4.9% in the U.K. in 2007, according to ONS [2009a].

REFERENCES

- Arbeleche-Gréla, S., and M.A.H. Dempster. "Econometric Modelling for Global Asset Management." Working Paper WP 2005-15, Judge Business School, University of Cambridge, 2005.
- Babel, B., E. Bornsdorf, and J. Kahlenberg. "Future Mortality Improvements in the G8 Countries." *Life & Pensions* (February 2008), pp. 42-48.
- Cairns, A.J.G. "Pension-Fund Mathematics." In *Encyclopedia of Actuarial Science*. Chichester, U.K.: John Wiley & Sons, 2004.
- Cambridge Systems Associates Limited (CSA). *Stochastics™ Manual*, 2008. Available at <http://www.cambridge-systems.com/solutions.html>.
- Clark, G.L., A.H. Munnel, and J.M. Orzag. *The Oxford Handbook of Pensions and Retirement Income*. Oxford, U.K.: Oxford University Press, 2006.
- Dempster, M.A.H. "Sequential Importance Sampling Algorithms for Dynamic Stochastic Programming." *Journal of Mathematical Sciences*, Vol. 133, No. 4 (2006), pp. 1422-1444.
- Dempster, M.A.H., I.V. Evstigneev, and K.R. Schenk-Hoppé. "Volatility-Induced Financial Growth." *Quantitative Finance*, Vol. 7, No. 2 (2007), pp. 151-160.
- . "Financial Markets: The Joy of Volatility." *Quantitative Finance*, Vol. 8, No. 1 (2008), pp. 1-3.
- Dempster, M.A.H., M. Germano, E.A. Medova, and M. Villaverde. "Global Asset-Liability Management." *British Actuarial Journal*, 9 (2003), pp. 137-216.
- Dempster, M.A.H., M. Germano, E.A. Medova, M.I. Rietbergen, F. Sandrini, and M. Scrowston. "Managing Guarantees." *Journal of Portfolio Management*, Vol. 32, No. 2 (2006), pp. 51-61.
- . "Designing Minimum Guarantee Funds." *Quantitative Finance*, Vol. 7, No. 2 (2007), pp. 245-256.
- Deutsche Bank. "U.K. Pensions: Regime Change ... Again." Research Report, Capital Structure Analysis Europe, 2005.
- Goldman Sachs. Presentation at Alpha Invest in Stockholm, 29-30 January 2008.

Medova, E.A., J.K. Murphy, A.P. Owen, and K. Rehman. "Individual Asset-Liability Management." *Quantitative Finance*, Vol. 8, No. 6 (2008), pp. 547-560.

Milevsky, M.A. *The Calculus of Retirement Income*. Cambridge, U.K.: Cambridge University Press, 2006.

Mulvey, J.M. "Generating Scenarios for the Towers Perrin Investment System." *Interfaces*, Vol. 26, No. 2 (1996), pp. 1-15.

Office for National Statistics (ONS). "Pension Contributions." Chapter 8 in *ONS Pension Trends*, 2009a. Available at http://www.statistics.gov.uk/downloads/theme_compendia/...pensiontrends/Pension_Trends_ch08.pdf.

———. "Pension Scheme Funding and Investment." Chapter 9 in *ONS Pension Trends*, 2009b. Available at http://www.statistics.gov.uk/...downloads/theme_compendia/pensiontrends/Pension_Trends_ch09.pdf.

Sandrini, F. "Simulators for the Risk Controlling System and Profitability Analysis for the German Retirement Product." *Quantitative Research Note*, Pioneer Investments, May 1, 2008.

Waring, M.B., and L.B. Siegel. "Don't Kill the Golden Goose! Saving Pension Plans." *Financial Analysts Journal*, Vol. 63, No. 1 (2007), pp. 31-45.

To order reprints of this article, please contact Dewey Palmieri at dpalmieri@ijjournals.com or 212-224-3675.