

Hedging Pension Plan Funding Ratio

Jon Palin
Cliff Speed

Presented at The Great Controversy: Current Pension Actuarial Practice in Light of
Financial Economics Symposium
Sponsored by the Society of Actuaries

Vancouver

June 2003

Copyright 2004 by the Society of Actuaries.

All rights reserved by the Society of Actuaries. Permission is granted to make brief excerpts for a published review. Permission is also granted to make limited numbers of copies of items in this monograph for personal, internal, classroom or other instructional use, on condition that the foregoing copyright notice is used so as to give reasonable notice of the Society's copyright. This consent for free limited copying without prior consent of the Society does not extend to making copies for general distribution, for advertising or promotional purposes, for inclusion in new collective works or for resale.

(note: The quotations appearing in this monograph are exact, except where capitalization and punctuation were changed in keeping with modern style and grammar guidelines.)

Abstract

This paper explores practical methods to hedge the fair value of pension scheme liabilities against changes in market conditions. The bond like nature of defined benefit pension liabilities enables hedge portfolios to be constructed. We consider how to find hedge portfolios for simple liabilities that are either fixed in monetary terms or linked to the consumer price index and also more complex benefits containing options. After discussing the difficulties of hedging we propose a new algorithm and apply it to practical problems.

1. Introduction

Both traditional actuarial logic and financial economics provide cases for hedging some or all of the liabilities of a defined benefit (DB) pension plan. This paper considers how a portfolio could be constructed to best hedge the liabilities.

Although the case for hedging liabilities has been long established and it is accepted that the hedge consists of bonds it is far from simple to construct a suitable portfolio. Legislative requirements in the United Kingdom make hedging particularly difficult. We discuss some of the difficulties in hedging complex liabilities and current approaches to the problem before proposing a new method.

Section 2 considers the case for hedging, discusses which liabilities should be hedged, and described the benefits we consider. Section 3 considers the close link between valuation and hedging, and Section 4 discusses the practical issues that make hedging difficult. Sections 5 and 6 consider current approaches to hedging and introduce our proposed method. Section 7 contains results from calculations using our approach, and Section 8 discusses these results and concludes.

Appendices provide supplementary information covering (a) state price deflators, (b) a difficult hedging problem, (c) our mathematical algorithm, and (d) the use of control variates.

2. Hedging

This section considers why we might want to hedge liabilities, which liabilities to hedge and the benefits to be considered.

2.1 Motivation for Hedging

From a financial economics point of view the wish to hedge the pension liabilities arises from the work of Sharpe (1976) and Black (1980) who set out the advantages to the plan sponsor of doing so. They argue that, although the sponsor could invest in assets that are expected to outperform the liabilities, this does not create value. This view arises from considering the investments of the pension fund in conjunction with the overall financing of the sponsoring company.

A company with risky pension assets could switch to assets that hedge the pension liabilities and restructure its balance sheet by issuing debt and buying back stock. If the balance sheet restructuring is carefully chosen then such action makes no difference to shareholder value to first order and produces second order gains by saving tax, reducing management costs and reducing discretionary benefit improvements. Hedging the liabilities also provides security for the pension fund members.

A notable practical example of this is The Boots Company, which moved its final salary plan entirely into bonds (Ralfe 2002).

We accept these views (discussed in more detail in Ralfe, Speed and Palin 2003), and they provide our motivation for investigating hedging portfolios. We believe our work is also of interest to those who wish to hedge liabilities for other reasons—perhaps the traditional actuarial logic that pensions that are close to payment should be hedged. Even if the plan is not hedged it is of interest to know the minimum risk-hedging portfolio and to assess asset performance against this scheme specific benchmark.

2.2 What to Hedge

We believe that the sponsor should seek to hedge the accrued pension liabilities, allowing for future indexation applicable to a leaver. In particular it should not allow for anticipated salary increases—future salary increases are a future liability.

2.3 Description of Benefits

Hedging pension benefits would be relatively easy if the pensions did not increase. However legislation makes U.K. pension benefits difficult to hedge. For example the 1995 Pensions Act requires limited price indexation (LPI) to be applied in payment to pensions accrued since April 6, 1997. The annual LPI increase is the annual change in the Retail Prices Index (RPI) subject to a minimum of 0 percent and a maximum of 5 percent.

Pensions must also be indexed before retirement. The statutory minimum increase is also described as LPI, but in this case the increase is not annual. The pension increase is the change in RPI between leaving and retirement subject to a minimum of 0 percent and a maximum of 5 percent per annum compounded over the whole deferment period.

3. Valuation

We next consider how to value DB pensions. There is a close link between valuation and hedging, since an asset portfolio that perfectly hedges liabilities will have the same value as the liabilities.

3.1 Basic Valuation Method

Our valuation method is based on this idea of no-arbitrage—two assets producing the same cash flows in all scenarios should have the same value. Consequently to value a non-increasing pension we could find a portfolio of bonds producing matching cash flows and take the value of that bond portfolio. An equivalent more convenient method is to discount the fixed cash flows using fixed-interest government bond yields of appropriate durations.

This basic approach works for non-increasing pensions, and for pensions directly linked to RPI we can use a similar method but using yields on index-linked bonds. However where the link to RPI is more complex, such as LPI increases, a more sophisticated method is required.

3.2 Valuation Method for Complex Increases

Liability definitions that include maximum and minimum pension increases can be thought of as options, where the underlying is the index to which increases are linked and the maximum and minimum pension increases correspond to strike prices of call and put options on the underlying. Because of the complexity of the options, in particular the annual compounding of LPI, we do not use an analytic approach.

Instead we use a stochastic model that is capable of producing inflation series and state-price deflators. Deflators allow the valuation of complex cash flows in a market-consistent manner. (For those unfamiliar with deflators we provide a brief description in Appendix A.) We use the inflation series to calculate the pension cash flows and then apply the deflators to the pension cash flows to calculate the capital value.

4. Difficulties of Hedging

Having decided to hedge liabilities, this section considers some of the difficulties we will face. Given simple liability cash flows, such as fixed cash flows, and a wide range of hedging assets, such as zero-coupon bonds of all terms, hedging is easy. We can simply select the right zero-coupon bonds to reproduce the fixed liability payments. In practice (at least) three difficulties arise. We describe these below.

4.1 Availability of Assets

At the time of writing the longest fixed-interest and index-linked U.K. government bonds (known as “gilts”) have redemption dates of 2036 and 2035 respectively. Pension liabilities are payable well beyond these dates. In addition gilts are not available for all redemption dates—there are just 11 index-linked gilts. Some fixed-interest gilts are “strippable”—the coupon and redemption payments can be traded separately, in effect creating zero-coupon gilts.

We note that swaps can be used to improve the hedge, and we are aware of a number of pension funds that hold swaps for this purpose, but we exclude swaps from this paper.

4.2 Presence of Maximum and Minimum Increases

The presence of maximum and minimum increases in the definition of LPI adds complexity. However it is possible to apply option-pricing theory to value and hedge these embedded options. Our results section considers how maximum and minimum increases affect the hedge portfolio.

4.3 Different Preretirement and Postretirement Increases

A further difficulty is caused by pensions that have differently defined increases before and after retirement. Consider a simple example, the “inflation-nil” case, which we shall return to later:

A lump sum is payable in two years’ time, and the amount payable is increased in line with inflation for the first year but receives no increase in the second year. That is, writing RPI_t for the RPI inflation index at time t (years), there is a single payment of RPI_1/RPI_0 after two years.

Even if we make available one- and two-year zero-coupon bonds, both fixed and index-linked, it is not possible to hedge this liability perfectly. (See Appendix B for an explanation of why a perfect hedge is not possible.) Our results section considers how we might hedge liabilities in this difficult case.

5. Current Approaches to Hedging

Before introducing our proposed approach to hedging we consider existing methods. These are listed in three parts, but there is considerable overlap between them.

5.1 Duration

Redington (1952) provides a good first approach to hedging liabilities. He considers the value of assets and liabilities to be a function of a single interest rate (no term structure) and derives conditions to be satisfied in order that asset and liability values should move together for small changes in this interest rate. For greater realism Redington's approach can easily be generalized to allow for parallel movements in interest rate term structures, but this permits arbitrage, as does his original method.

Panjer (1998) discusses Redington's (1952) work and similar methods to quantify the interest rate sensitivity of securities.

5.2 Option Pricing

We said previously that the presence of maximum and minimum pension increases is a form of option. A modification of the Black-Scholes formula (Black and Scholes 1973) can provide a reasonable approximation to the hedging portfolio for an LPI increase over one year¹ but will not capture the full complexity, particularly the presence of different increases before and after retirement.

Other techniques from option pricing can also be used such as calculating the sensitivity of assets and liabilities to changes in market conditions (the "Greeks"). Hull (2002) describes a range of such techniques.

5.3 Analytical Approach

Van Bezooeyen, Exley, and Smith (1997) and Huang and Cairns (2002) have taken an analytical approach, choosing a functional form for interest rates and inflation and solving to find hedge portfolios.

6. Proposed Method

¹ We are aware of a pension consultancy that uses this approach.

This section introduces our proposed method for hedging. Before describing our method we discuss our reasons for choosing it.

6.1 Motivation

We recognize that, in most cases, a perfect hedge won't be possible—the “inflation-nil” case discussed earlier shows this. Consequently we need a way to compare the effectiveness of different possible hedges and choose between them, while accepting that there is no single correct measure.

We want to be market consistent, in particular considering full term structures of nominal and real interest rates, and we calibrate these to market data. This is important in ensuring that the portfolio produced by our method can actually be bought for the price given by our method.

We want our method to be flexible—to cope with an arbitrarily complex asset model and be able to address other hedging problems.

6.2 Method

In deciding the effectiveness of a hedge we choose to measure the expected squared difference between the asset and liability values. We recognize that other measures have merit, for example, placing more emphasis on deficits than surpluses, but ignore these in this paper².

In practice there may not be a great deal of difference in the portfolio produced using different measures. Since we have an arbitrage free model and the asset and liability values are initially equal, a strategy that made the deficit zero would necessarily make the surplus zero. So we would not expect a portfolio arising from an attempt to reduce a function of the deficit to be radically different to a portfolio arising from an attempt to reduce the same function of the difference between asset and liability values.

Using this measure we choose the asset portfolio so that:

- Asset value now equals liability value at time 0 (now).
- Expected squared difference between asset and liability values at time t (a short time later) is minimized.

² The mathematical treatment contains a parameter W , which could be used to place more weight on surplus or deficit. However we choose to use $W = 1$ in this paper, placing equal weight on surpluses and deficits.

In this paper we use annual time steps ($t = 1$), so “a short time later” means one year, but the method can use other periods. A time step of one year may seem long compared to other financial algorithms, but in practice few pension schemes review their investments on a more regular basis.

We use deflators as described earlier to get the liability values—this ensures market-consistency and enables us to cope with difficult liability definitions. We also choose to use deflators to obtain the asset values. The liability values calculated using deflators will not be exact because of sampling error, and calculating asset values using deflators appears to help because the sampling errors of assets and liabilities are related. We have tried using the yields from our stochastic model to calculate asset values, but we have found that using deflators produces better results for the hedging portfolio.

The problem as stated previously is an exercise in constrained optimization—minimizing expected squared difference between asset and liability values after a year, subject to equality of asset and liability values now. We solve this problem using Lagrangian multipliers and after some manipulation obtain a solution for the portfolio in terms of deflators and asset and liability cash flows. Mathematical details are provided in Appendix C.

To calculate the portfolio, we (1) use our suitably calibrated asset model to generate deflators and inflation for many simulations, (2) derive asset and liability cash flows from the simulated inflation and (3) calculate the portfolio from these cash flows and the deflators.

The next section discusses the results we obtain by applying this method.

7. Results

7.1 Assumptions

Our results below make use of a proprietary stochastic asset model. This is an arbitrage free model calibrated to full term structures of nominal and real interest rates. All results in this paper are based on a calibration to U.K. yield curves at Jan. 1, 2003.

It is worth noting that, while the initial yield curves are calibrated to market conditions, there is no way to derive a market parameter for the volatility of inflation. This parameter requires a subjective estimate and will have a critical impact on our numerical results.

The assets available for the algorithm to select are zero-coupon government bonds, both fixed and index-linked, of all integer terms. (These gilts are hypothetical and are not issued by the U.K. government—although a reasonable approximation is possible for the fixed-interest part using strippable gilts.)

We could use realistic assumptions for mortality and other demographics, but this isn't a difficult part of the problem and could obscure more useful findings. Consequently we look at a fixed preretirement period and a fixed postretirement period.

7.2 Experiments Done

Our calculations can be split into four types:

1. Cases where a perfect hedge is possible and the answer can be easily calculated without the model. We use these as a basic check of our approach.
2. Cases where a minimum and maximum pension increase applies.
3. Cases where the increase is different in deferment and retirement, including the “inflation-nil” case.
4. A realistic case.

7.3 Basic Checks

In this section we consider cases where there is a portfolio that hedges the liabilities perfectly and we can calculate this readily.

Our first basic checks consisted of cases where the pensions did not increase in deferment or in payment. In this case it is easy to see that the hedging assets are fixed zero-coupon gilts redeeming at the times of the pension payments, with nominal equal to the pension payment. Our algorithm produced portfolios consisting of precisely these assets.

We then repeated these calculations for cases where the pensions increase in line with full RPI in deferment and in payment. In this case the hedging assets are index-linked zero-coupon gilts and our algorithm produced the correct portfolios.

7.4 Minimum and Maximum Pension Increases

The first difficult case where the algorithm is of interest is the case where minimum and maximum pension increases apply. We first consider pensions in payment for one year that increase with the RPI index subject to a minimum of 0 percent each

year and some maximum increase, and then see how the portfolio changes as the maximum is varied.

These results are summarized in Table 1. We show the proportion of assets by initial value invested in fixed-interest and index-linked gilts, and the duration of the portfolio.

Table 1
Results for Pensions Payable for One Year, With a Zero-Percent Minimum Increase and Different Maximum Increases

Pension Type	Fixed-Interest	Index-Linked	Duration
min 0% max 0%	100.0%	0.0%	1.000
min 0% max 1%	99.3	0.7	1.000
min 0% max 2%	88.3	11.7	1.000
min 0% max 3%	39.4	60.6	1.000
min 0% max 4%	2.6	97.4	1.000
min 0% max 5%	0.0	100.0	1.000

Although we made available zero-coupon gilts of terms from one to five years the duration of the assets matches the duration of the liabilities in all cases. This is satisfying, particularly as our algorithm has no concept of duration and so is not constrained to produce the matching duration.

We see that, as we raise the maximum pension increase, the proportion of assets invested in index-linked gilts is also raised. This makes sense. With an increasing maximum increase the pension moves from being a fixed pension to being more like a pure index-linked pension.

The numerical results are sensitive to the inflation model we use. In particular as we are only projecting inflation over one year from the calculation date the inflation series has low volatility. The expected inflation (derived from the market) is 2.90 percent and our model gives inflation outside the range 0 percent to 5 percent in only nine of 50,000 simulations. Consequently an LPI pension with a minimum increase of 0 percent and a maximum increase of 5 percent is almost identical to an RPI pension, and the hedging asset is the index-linked gilt.

We next consider the portfolio to hedge an LPI pension payable over five years. Results are shown in Table 2. The values shown are percentages of the initial value of the portfolio.

Table 2
Suggested Portfolio to Hedge an LPI Pension (Zero-Percent Minimum Increase and 5 Percent Maximum Increase Applied Annually) Payable for Five Years

Asset Type and Term	Term 1	Term 2	Term 3	Term 4	Term 5	Total
Index-Linked Gilt	43.3%	18.5%	14.5%	14.3%	11.1%	101.7%
Fixed-Interest Gilt	-22.7	1.9	5.6	5.4	8.2	-1.7
Total	20.6	20.4	20.0	19.7	19.3	100.0

There are three points of interest in the results. First, the total gilt holding of terms 1 to 5 are approximately equal at around 20 percent each, but less is held for longer terms as these payments are less valuable (since the LPI pension increase is less than the discount rate). This fits with our expectation and provides some reassurance in the method.

Second, we see that the proportion invested in fixed-interest gilts increases as the term increases. This is again reasonable. The payments at time 5 depend on inflation over five years and this is much less predictable than inflation over the next year. For example expected inflation in year 5 is 2.14 percent (derived from the market) and our econometric model gives a 42 percent chance that the RPI increase will be below 0 percent or above 5 percent, so there is a 42 percent chance that either the maximum or minimum increase will apply in year 5. As the chance of a known 0 percent or 5 percent payment increases it makes sense to have an increased holding of fixed interest gilts to meet these known payments.

Finally, it is striking that there is a large short position in term-1 fixed-interest gilts (the closest asset to cash that we make available). Indeed this is sufficiently large that the net position in fixed-interest gilts is negative. In considering why this should be it is important to remember that the portfolio suggested is the portfolio to be held over the first year, and it will be rebalanced at time 1.

The portfolio suggested by the model is a compromise between providing exposure to inflation over the first year, as this will have a big impact on the liability value at time 1, and holding assets of the right duration to hedge against changes in yield curves. The portfolio produced has approximately a 100 percent exposure to the (fairly predictable) inflation over year 1 while maintaining a good duration hedge.

Both the short position in term-1 fixed-interest gilts and the small net holding in fixed-interest gilts are also seen for other pension terms. Table 3 shows summary results for LPI pensions payable over one to five years.

Table 3
Results for LPI Pensions (Zero-Percent Minimum Increase and 5 Percent Maximum Increase) Payable for a Varying Number of Years

Pension Type	Fixed-Interest	Index-Linked	Duration
One-Year LPI	0.0%	100.0%	1.000
Two-Year LPI	2.6	97.4	1.497
Three-Year LPI	2.7	97.3	1.991
Four-Year LPI	0.7	99.3	2.480
Five-Year LPI	-1.7	101.7	2.966

7.5 “Inflation-Nil” and Similar Cases

The next case of interest is the “inflation-nil” case mentioned earlier—a single payment made after two years, with an RPI increase in the first year and no increase in the second year. Although there is no minimum or maximum pension increase, the different type of increase in years 1 and 2 means that this pension cannot be hedged perfectly. Our model produces the results in Table 4.

Table 4
Results for the “Inflation-Nil” Case

Asset Type and Term	Term 1	Term 2
Index-Linked Gilt	+28.4%	+22.9%
Fixed-Interest Gilt	-28.4	+77.1

We also consider the “nil-inflation” case where there is an RPI increase in the second year but no increase in the first year. Results are shown in Table 5.

Table 5
Results for the “Nil-Inflation” Case

Asset Type and Term	Term 1	Term 2
Index-Linked Gilt	-27.6%	+76.9%
Fixed-Interest Gilt	+27.5	+23.1

The portfolios suggested for both the “inflation-nil” and “nil-inflation” cases have duration of 2.000, the same as the liability duration. Again our algorithm produces a duration hedge even though this is not a constraint of the method.

The “inflation-nil” and “nil-inflation” cases are defined in a similar way, but with the order of the increases changed. The portfolios produced are also similar with the roles of index-linked and fixed-interest gilts interchanged. We choose to discuss the portfolio for the “nil-inflation” case.

It is important to recall that our method produces a portfolio to be held for the first year, and this will be rebalanced at that time. Once we get to time 1 we want to be holding just the term-2 index-linked gilt (which will then have a term of 1) in order to have exposure to inflation over the second year. However holding this from the start also gives us exposure to inflation over the first year. The short position in the term-1 index-linked gilt removes some of this exposure.

As explained earlier there is no perfect solution to the “inflation-nil” or “nil-inflation” problems. We should not expect to be able to explain the portfolios perfectly. Also, the portfolios produced are very sensitive to the inflation model we use. The model used for the previous results assumes that inflation is more predictable over the short term than over the longer term. Table 6 shows the result for the “inflation-nil” case using an alternative inflation model in which the volatility of inflation does not depend on the time.

Table 6
Results for the “Inflation-Nil” Case Using an Alternative Inflation Model

Asset Type and Term	Term 1	Term 2
Index-Linked Gilt	+82.2%	+16.1%
Fixed-Interest Gilt	-82.1	+83.7

While the broad pattern of the portfolio is the same, the holdings of term-1 gilts are much more extreme in Table 6. This shows that the choice of inflation model has a large impact on the results.

For completeness we show in Table 7 the portfolio for the “nil-inflation” case using the alternative inflation model in which inflation volatility does not vary over time. Again the change of model has a big impact on the portfolio.

Table 7
Results for the “Nil-Inflation” Case Using an Alternative Inflation Model.

Asset Type and Term	Term 1	Term 2
Index-Linked Gilt	-82.5%	+83.9%
Fixed-Interest Gilt	+82.4	+16.4

7.6 A More Realistic Case

We finally consider a more realistic case where the pension has a deferred period of five years and a payment period of five years. The pension is increased with LPI in deferment and in payment³.

Table 8
Results for a More Realistic Pension—Five Years in Deferment, Five Years in Payment, With LPI Increases

Asset Type and Term	Term 1	Term 2	Term 3	Term 4	Term 5
Index-Linked Gilt	-1.7%	+1.3%	-0.1%	+2.5%	+9.6%
Fixed-Interest Gilt	+1.8	-1.4	+0.0	-2.1	-9.8
Total	+0.1	-0.1	-0.1	+0.4	-0.1

Asset Type and Term	Term 6	Term 7	Term 8	Term 9	Term 10	Total
Index-Linked Gilt	+17.9%	+10.7%	+12.9%	+8.9%	+10.1%	72.2%
Fixed-Interest Gilt	+2.9	+9.7	+7.1	+10.5	+8.9	27.8
Total	+20.8	+20.5	+20.0	+19.5	+19.1	100.0

The majority of the portfolio is invested in index-linked gilts. This is reasonable as the pension is largely index-linked. There is a low chance that inflation will be above 5 percent compound or below 0 percent compound during the deferred period, so the maximum and minimum increases are unlikely to have a large impact. They have more of an impact in payment when the limits are applied annually but we would still expect fixed-interest gilts to be in the minority.

We see that, from terms 6 to 10, when the pension is payable, there is a fairly even spread of total assets at around 20 percent of the liability value, and that the total investment decreases from term 6 to term 10 as the value of the payments decrease. There are much smaller positions from terms 1 to 5 during the deferred period. It is

³ As described earlier, LPI in deferment is applied over the whole deferred period—not year by year.

hard to know whether the holdings of gilts of terms 1 to 4 are fulfilling a useful purpose or merely an artifact due to lack of convergence of the stochastic simulations.

The holding of term-5 gilts does seem to be a genuine result. Repeating the calculations for deferred periods of 3 or 4 also shows a significant long index-linked and short fixed-interest position in the term immediately before the pension comes into payment.

8. Discussion and Conclusion

Hedging pension liabilities is difficult, more difficult than we thought when we started to write this paper. We don't claim to have found a full solution to the problem; this paper represents work in progress. However we believe that our method produces reasonable, sensible results for the cases described previously, even if we can't explain the answers fully.

The identification of the "inflation-nil" case is useful in explaining why the hedging problem is difficult. We hope that this, along with our "realistic case," should serve as a useful benchmark for comparing different approaches to hedging.

The main advantage of our method is its flexibility. It does not depend on a particular model of inflation or yield curves, and any arbitrage-free econometric model could be used. However the need for stochastic simulations causes problems. For liabilities with short terms it is feasible to make a large number of simulations and get good convergence to a suitable portfolio. But the computation required grows as the term of the liabilities increases and even in the case we have considered it is not clear whether some parts of the portfolio are genuinely useful in hedging or just spurious results through lack of convergence. We are considering the use of variance control techniques (in particular control variates) to reduce the effect of this problem. We describe this in Appendix D.

A further concern highlighted by the results is the sensitivity to the choice of inflation model. Although we can calibrate the initial yield curves to market conditions and, hence, use a market expectation of inflation, we need to assume a process and volatility structure for inflation. This is difficult to do, particularly as changes in recent monetary policy⁴ make it difficult to use historic data. The results in Tables 4 and 6 showed the large impact on the hedging portfolio from changing

⁴ For example, in 1997, the Chancellor of the Exchequer in the United Kingdom, Gordon Brown, introduced an inflation target of 2.5 percent per annum and passed responsibility for setting interest rates and meeting this target to the Bank of England. This action had an immediate impact on market expectations of inflation and also affected the likely volatility of inflation. Similar targets operate in other countries.

the volatility structure of inflation and however sophisticated our models the hedge portfolio will depend on the assumed inflation process.

We have a long list of further work to do. In particular our portfolios are those to be held for the first year and we have not considered to what extent they would need to be rebalanced after one year. The cost of such rebalancing may be large relative to the liability values and may suggest that portfolios should be rebalanced more or less frequently than annually. We also wish to test our method against historic market conditions.

Acknowledgments

The mathematical algorithm used in this paper has been developed from an unpublished idea by Andrew Smith. We are grateful to Andrew and also to Tim Gordon who drew Andrew's work to our attention. We also wish to thank Ítalo Rocha Souza and Miles Blackford for running the stochastic model and assisting with our computer programs and Stuart Jarvis for reviewing (and correcting) our algorithm.

Appendix A Deflators

In this section, we give a very brief introduction to “deflators.” For more information, we refer the interested reader to Jarvis, Southall and Varnell (2001).

We would like to be able to discount the cash flows from an asset to find the value of that asset. This causes problems if we use a deterministic discount factor as we need a different discount rate for each asset to allow for the risk of the cash flows; for example, cash flows from a risky equity should be discounted at a higher rate than cash flows from a safer bond.

A deflator can be thought of as a “stochastic discount rate” that can be used for all assets. A deflator is a time-dependent random variable and has the property that

If asset i produces cash flows C_t^i (random variable) and the deflator is D_t (random variable) then $E(\sum_t D_t C_t^i)$ gives the value of asset i .

The same approach also allows other assets and liabilities to be priced if their cash flows depend on the assets to which the deflator has been calibrated; for example, if our model has a deflator that will give the market value for an equity and a gilt, then it can also give the market-consistent price for an outperformance option based on those assets.

The presence of deflators taking only positive values ensures that an econometric model is arbitrage free. Conversely if an econometric model is arbitrage free then it is possible (although often difficult) to build deflators into the model.

Appendix B
Why the “Inflation-Nil” Case Cannot Be Hedged Perfectly

In this section, we write $r_{a,b}$ for the real interest rate and $n_{a,b}$ for the nominal interest rate with remaining term b at time a (both rates being continuously compounded), and RPI_t for the Retail Prices Index at time t . We assume $RPI_0 = 1$. Table B1 shows the values of our four assets at time 0 and immediately before time 1.

Table B1

Asset	Value at Time 0	Value Just Before Time 1
Fixed-Interest Term 1	$\exp(-n_{0,1})$	1
Fixed-Interest Term 2	$\exp(-2n_{0,2})$	$\exp(-n_{1,1})$
Index-Linked Term 1	$\exp(-r_{0,1})$	RPI_1
Index-Linked Term 2	$\exp(-2r_{0,2})$	$RPI_1 \exp(-r_{1,1})$

The liability value immediately before time 1 is $RPI_1 \exp(-n_{1,1})$. If the liability could be hedged perfectly then we could set up a portfolio of the four assets at time 0, which would have the same value at time 1 as the liability. This is equivalent to finding four coefficients so that the liability value is equal to a linear combination of the asset values for all values of $r_{1,1}$, $n_{1,1}$ and RPI_1 . Equivalently if the liability could be hedged perfectly then we could find constants a, b, c, d so that:

$$a.1 + b.\exp(-n_{1,1}) + c.RPI_1 + d.RPI_1 \exp(-r_{1,1}) = RPI_1 \exp(-n_{1,1}) \text{ for all } r_{1,1}, n_{1,1} \text{ and } RPI_1.$$

So we have to solve infinitely many equations in four unknowns. It is easy to find five sets of values for $r_{1,1}$, $n_{1,1}$ and RPI_1 so that the system of equations is not solvable, showing that the “inflation-nil” case cannot be hedged.

Appendix C Mathematical Algorithm

C1. Assumptions

We assume that cash flows take place at integer times $t = 1, 2, \dots, T$.

C2. Notation

Write:

- δ for a small time period.
- D_t for the deflator at time t (a random variable).
- $C_t^{Asset_i}$ for the cash flow from asset i ($i = 1, 2, \dots, N_{assets}$) at time t (a random variable).
- $C_t^{Liability}$ for the liability cash flow at time t (a random variable).
- p_i for the (nominal) amount invested in asset i . The vector \underline{p} is the portfolio we want to find.
- W for a stochastic weighting factor (a random variable).

C3. Basic Algorithm

We calculate the prices of an asset from its cash flows using deflators. We have that the total asset value at time $0 + \delta$, immediately after a cash flow is

$$\frac{1}{D_0} E_0 \left[\sum_{t=1}^T D_t \cdot \left(\sum_{i=1}^N C_t^{Asset_i} \cdot p_i \right) \right]$$

the total asset value at time $1 - \delta$, immediately before a cash flow is

$$\frac{1}{D_1} E_1 \left[\sum_{t=1}^T D_t \cdot \left(\sum_{i=1}^N C_t^{Asset_i} \cdot p_i \right) \right]$$

Similarly, the liability value at time $0 + \delta$ is

$$\frac{1}{D_0} E_0 \left[\sum_{t=1}^T D_t \cdot C_t^{Liability} \right]$$

and the liability value at time $1 - \delta$ is

$$\frac{1}{D_1} E_1 \left[\sum_{t=1}^T D_t \cdot C_t^{Liability} \right].$$

We aim to minimize the (stochastically weighted) expected squared difference between asset and liability values at time $1 - \delta$, subject to the asset and liability values at time $0 + \delta$ being equal.

That is, we want to minimize

$$E \left[W \cdot \left(\frac{1}{D_1} E_1 \left[\sum_{t=1}^T D_t \cdot C_t^{Liability} \right] - \frac{1}{D_1} E_1 \left[\sum_{t=1}^T D_t \cdot \left(\sum_{i=1}^N C_t^{Asset_i} \cdot p_i \right) \right] \right)^2 \right] \quad (1)$$

subject to

$$\frac{1}{D_0} E_0 \left[\sum_{t=1}^T D_t \cdot \left(\sum_{i=1}^N C_t^{Asset_i} \cdot p_i \right) \right] = \frac{1}{D_0} E_0 \left[\sum_{t=1}^T D_t \cdot C_t^{Liability} \right] \quad (2)$$

We set the Lagrangian function to be a combination of equations (1) and (2) that is,

$$Z = E \left[W \cdot \left(\frac{1}{D_1} E_1 \left[\sum_{t=1}^T D_t \cdot C_t^{Liability} \right] - \frac{1}{D_1} E_1 \left[\sum_{t=1}^T D_t \cdot \left(\sum_{i=1}^N C_t^{Asset_i} \cdot p_i \right) \right] \right)^2 \right] + \lambda E_0 \left[\sum_{t=1}^T D_t \cdot \left(C_t^{Liability} - \sum_{i=1}^N C_t^{Asset_i} \cdot p_i \right) \right]$$

And then solve for \underline{p} so that $\frac{\partial Z}{\partial p_i} = 0$ for all i and $\frac{\partial Z}{\partial \lambda} = 0$.

After some manipulation and application of the tower law for conditional expectation, we derive the equations:

$$\underline{G} \cdot \underline{p} + \frac{1}{2} \lambda \underline{a} = \underline{h} \quad \text{and} \quad \underline{a}^T \cdot \underline{p} = l, \quad \text{respectively,}$$

where \underline{G} is a matrix, \underline{h} and \underline{a} are vectors and l is a scalar, with

$$g_{ij} = E \left[\frac{W}{D_1^2} \sum_u \sum_t D_u D_t C_u^{Asset_i} C_t^{Asset_j} \right],$$

$$h_i = E \left[\frac{W}{D_1^2} \sum_u \sum_t D_u D_t C_u^{Asset_i} C_t^{Liability} \right],$$

$$a_i = E \left[\sum_t D_t C_t^{Asset_i} \right] \quad \text{the value of the assets, and}$$

$$l = E \left[\sum_t D_t C_t^{Liability} \right] \quad \text{the value of the liabilities.}$$

These equations can be solved by calculating

$$\lambda = 2 \frac{\underline{a}^T \underline{\underline{G}}^{-1} \underline{h} - l}{\underline{a}^T \underline{\underline{G}}^{-1} \underline{a}}$$

$$\underline{p} = \underline{\underline{G}}^{-1} (\underline{h} - \frac{1}{2} \lambda \underline{a})$$

Note: In our results we have used the value $W \equiv 1$.

Appendix D Control Variates

We have not used control variates in our paper but believe that the technique may offer improved convergence to our stochastic simulations.

D1. Method

Suppose we want to find the mean of the random variable X . If we can only simulate X itself then our best estimate is just $E[X]$. Suppose however that we can also simulate the N random variables Z_i using the same simulations, and that the values $E[Z_i]$ are known for all $i = 1, 2, \dots, N$.

In this case we construct a new random variable Y

$$Y = X + \sum_{i=1}^N \beta_i (Z_i - E[Z_i]),$$

where β_i are constants.

Now $E[Y] = E[X]$ so Y is an unbiased estimator for X . We choose β_i so that $Var(Y)$ is minimized.

$$V = Var(Y) = Var(X) + 2 \sum_{i=1}^N \beta_i Cov(X, Z_i) + \sum_{i=1}^N \sum_{j=1}^N \beta_i \beta_j Cov(Z_i, Z_j)$$

$$\frac{\partial V}{\partial \beta_i} = 2Cov(X, Z_i) + 2 \sum_{j=1}^N \beta_j Cov(Z_i, Z_j)$$

To minimize V we want to set all $\frac{\partial V}{\partial \beta_i}$ to zero. Define a matrix $\underline{\underline{C}} = \{c_{ij}\}$ and vectors

$\underline{b} = \{\beta_i\}$ and $\underline{x} = \{x_i\}$, where $c_{ij} = Cov(Z_i, Z_j)$ and $x_i = Cov(X, Z_i)$.

Then we want to solve $\underline{\underline{C}}\underline{b} = -\underline{x}$, so the values of β_i are given by $\underline{b} = -\underline{\underline{C}}^{-1}\underline{x}$.

D2. Application

Our algorithm described previously consists of calculating four expectations and using these to find the hedging portfolio. Using control variates to reduce the sample error of these expectations may produce a better hedging portfolio.

References

- Black, F. 1980. The tax consequences of long-run pension policy, *Financial Analysts Journal* 36: 627–38.
- Black, F., and M.S. Scholes. 1973. The pricing of options and corporate liabilities, *Journal of Political Economy* 81: 637–54.
- Huang, H-C., and A.J.G. Cairns. 2002. Valuation and hedging of LPI liabilities. Preprint. <http://www.ma.hw.ac.uk/~andrewc/papers/ajgc28.pdf>
- Hull, J. 2002. *Options, futures, and other derivatives*. Upper Saddle River, NJ: Prentice Hall.
- Jarvis, S., F. Southall, and E. Varnell. 2001. Modern valuation techniques Paper presented to the Staple Inn Actuarial Society. <http://www.sias.org.uk/papers/mvt.pdf>.
- Panjer, H.H. ed. 1998. *Financial economics: With applications to investments, insurance and pensions*. Schaumburg, IL: The Actuarial Foundation.
- Ralfe, J. 2002. Why move to bonds? *The Actuary* (March). http://www.the-actuary.org.uk/monthsissues_frames/articles//02_03_28-29.pdf.
- Ralfe, J., C.A. Speed, and J. Palin. 2003. Pensions and capital structure: Why hold equities in the pension fund? Paper presented to the Society of Actuaries' "The Great Controversy: Current Pension Actuarial Practice in Light of Financial Economics Symposium," June.
- Redington, F. 1952. Review of the principles of life-office valuations, *Journal of the Institute of Actuaries* 78: 286–340.
- Sharpe, W.F. 1976. Corporate pension funding policy, *Journal of Financial Economics* 183–93.
- van Bezoooyen, J.T.S., C.J. Exley, and A.D. Smith. 1997. A market-based approach to valuing LPI liabilities. Paper presented to the joint "Institute and Faculty of Actuaries Investment Conference." http://www.gemstudy.com/DefinedBenefitPensionsDownloads/Market_Based_Approach_to_Valuing_LPI_Liabilities.pdf.