Pricing the Risk on Longevity Bonds

Andrew J.G. Cairns[ab], David Blake[c], Paul Dawson[c], Kevin Dowd[d]

February 22, 2005

Abstract

The impact of increasing longevity on pension provision has become a major concern recently. More effective management of longer-term mortality risk has been made possible following a recent and innovative bond issue. We analyse the risk premium embedded in this instrument and argue that demand for further such instruments and innovations can be expected.

Introduction

Recent years have seen an increasing desire on the part of insurers, reinsurers and pension plans to hedge their mortality risks more effectively. December 2003 saw the issue of a 3-year bond by Swiss Re and Vita Capital. This groundbreaking issue was the first floating-rate bond to link the return of principal solely to a mortality index. More specifically, for Swiss Re, the bond was designed to help hedge their exposure to catastrophic mortality risks such as major epidemics or a terrorist attack (on a scale far greater than 9/11).

Buyers of the bond included a number of pension funds, for which the bond represented a hedge against a sudden fall in their pension liabilities resulting from significantly higher deaths in the short term than expected.[1] However, for primary insurers and pension plans, this bond was only a hedge against one particular form of extreme short-term mortality risk. These financial institutions are also exposed to many other forms of mortality risk as well, and the most important of these are longer-term mortality risks.

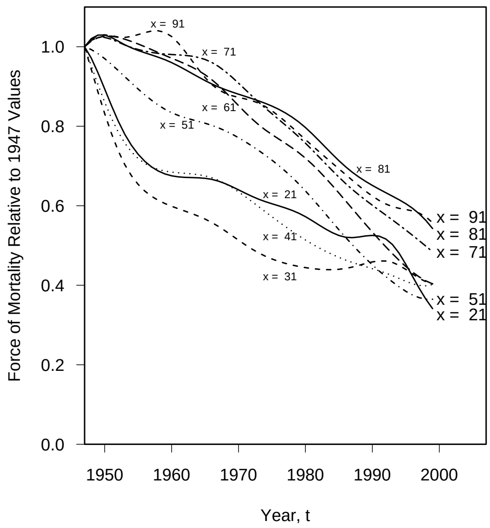

Figure 1: The evolution of mortality: Fitted values using P-splines for the instantaneous mortality rate, $\hat{\mu}(t, x)$, relative to the 1947 value for the years $t = 1947$ to 1999 and for ages $x = 21$, 31, 41, 51, 61, 71, 81 and 91.

It is now acknowledged that changes in mortality rates over time are only partially, rather than fully, predictable. Figure 1 shows how mortality rates at different ages have evolved over time relative to their values in 1947 for UK males, assured lives (data adapted from Currie *et al*, 2004). These data (the values show the smoothed instantaneous mortality rate[2]) allow us to make three observations. First, improvements over the the past 50 years have been significant. Second, these improvements look, to some extent, random. Third, the pattern of improvements over time has been different at different ages.

One year on from the issue of the Swiss Re bond, in November 2004, BNP Paribas announced that it had arranged for a longevity bond to be issued by the European Investment Bank (EIB) that goes a very long way towards providing a solution for financial institutions looking for instruments to hedge their long-term systematic mortality risks. The total value of the issue was £540 million, and was primarily intended for purchase by UK pension funds. The concept and usefulness of longevity bonds have been discussed for a number of years (Cox *et al.*, 2000,

[a]Actuarial Mathematics and Statistics, School of Mathematical and Computer Sciences, Heriot-Watt University, Edinburgh, EH14 4AS, United Kingdom.

[b]Corresponding author: E-mail A.Cairns@ma.hw.ac.uk

[c]Pensions Institute, Cass Business School, City University, 106 Bunhill Row, London, EC1Y 8TZ, United Kingdom.

[d]Centre for Risk & Insurance Studies, Nottingham University Business School, Jubilee Campus, Nottingham, NG8 1BB, United Kingdom.

[1]Such an event is, of course, beneficial financially for pension funds. Pension funds were, nevertheless, prepared to buy the bond because it reduced variability in the asset-liability ratio and because the bond offered an attractive return relative to conventional bonds.

[2]Actuaries usually refer to this as the *force of mortality*.

and Blake & Burrows, 2001). However, it has taken time for the capital markets to develop the finer implementation details of these contracts (even though here the detail is relatively simple), for both potential issuers and sufficient investors to decide that the time is right.

The structure of the bond is quite simple. Payments are linked to a cohort survivor index based on the realised mortality rates of English and Welsh males aged 65 in 2003. Payments on the bond are based on an initial annuity of £50 million. We, therefore, set $S(0) = £50$ million to be the base value for the index on the issue date. The payment at time t is then $S(t)$ where the values are updated recursively as follows

$$S(t+1) = S(t).\big(1 - q(t-l, 65+t)\big) \qquad (1)$$

where $q(t, x)$ is the crude (unsmoothed) population mortality rate for males aged x in year t for England and Wales. These rates are published on a regular basis by the Office for National Statistics[3] some time after the end of year t which explains the requirement for a time lag of l in (1).

The bond was priced by taking the projected survival rates produced by the UK Government Actuary's Department[4] (GAD). Projected coupons (payable annually) are plotted in Figure 2) and these were discounted at LIBOR minus 35 basis points to give the issue price.

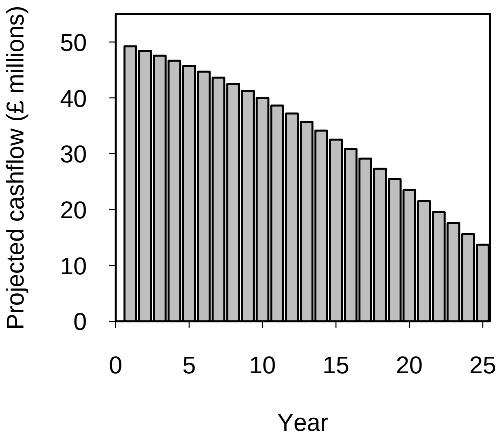

Figure 2: Projected coupons under the BNP Paribas longevity bond.

[3]See http://www.statistics.gov.uk
[4]See http://www.gad.gov.uk

Was it a good deal?

The fact that the bond is to be issued by the AAA-rated EIB should ensure that it attracts a strong rating. However, the risk associated with differences between experienced mortality and the GAD projection will be borne by AA-rated Partner Re. In order that the AA status of both Partner Re and of BNP Paribas do not adversely affect the longevity bond's rating, collateral arrangements have been made to minimise this risk.

One might wish to speculate whether the pension plans got a good deal or not? The background theory to the pricing of such securities can be found in Cairns, Blake and Dowd (2004) and in Lin and Cox (2004). Here, however, we will take a much simpler approach given that we know what the issue price is.

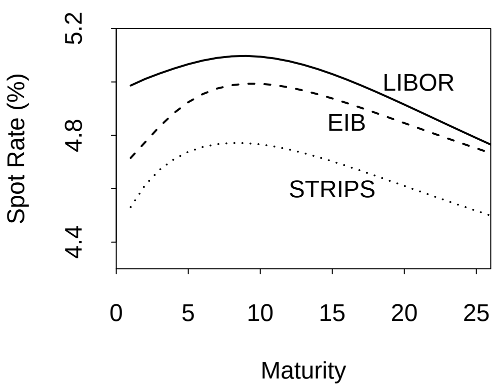

Figure 3: Annualised spot rates on 18 November 2004 for LIBOR (solid curve), the EIB (secondary market) (dashed) and gilt STRIPS (dotted).

Suppose, first of all, that actual mortality improvements are deterministic and match the GAD's projected mortality rates. The discount rate of LIBOR-35 contrasts with AAA-rated, fixed-interest EIB bonds which are normally funded at LIBOR-15 in the primary market.[5] We can see from Figure 3 that LIBOR-35 places us close to the gilts yield curve. The spread of 20 basis points between LIBOR-15 and LIBOR-35 comes mainly from the fact that the future development of mortality is stochastic rather than deterministic.[6]

[5]Note though, that the secondary market in fixed-interest EIB bonds implies pricing at a range of rates both above and below LIBOR-15, depending on term to maturity: see Figure 3.

[6]Lower investment management fees also contributes to the spread of 20 bps but to a much lesser extent. Specifically, the longevity bond is intended to be used as

As a consequence, even if the GAD projection is accurate *ex ante* as a *mean* trajectory, pension funds are paying a premium to reduce their exposure to this risk by investing in the bond. There are four main risks:

- Non-systematic (or sampling) mortality risk. Even if the survival probabilities are known the actual number who die each year is random.
- Systematic mortality improvement risk. For a given model and parameter set, future mortality rates will develop in a stochastic fashion.
- Parameter risk. Parameters are estimated imprecisely because of the limited amount of data available.
- Model risk. A number of models will fit the limited historical data reasonably well and we cannot tell which of these, if any, is 'correct'. Importantly, different models may give rise to different projections of the future.

Of these risk factors only the first is relatively insignificant because the chosen index has a sufficiently large underlying population. A specific aspect of model risk exists in the use of the GAD projection as the benchmark for pricing. If it is believed that this is significantly different from the true mean then the price paid needs to reflect this bias.

The other risk factors are more difficult to handle, but we can illustrate the issues involved with the following simple model for the development of mortality rates $q(t, x)$:

$$q(t,x) = \frac{\exp(a(t) + b(t)x)}{1 + \exp(a(t) + b(t)x)}$$

where $(a(t), b(t))$ is a vector-$ARIMA(1, 1, 0)$ time series.[7] Figure 4 shows confidence intervals for the proportion of the cohort surviving to different ages under this model.[8] Other models give similar median projections but confidence intervals of varying widths. However, the median projections for these different models all lie close to the GAD's projection in Figure 2.

The variance of the log of the proportion of survivors at age 90 in this graph is about 0.014.

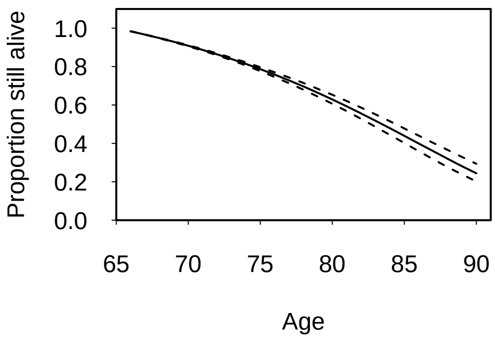

Figure 4: Median (solid line) and 95% confidence intervals (dashed lines) for future cashflows.

One can take this as an estimate of the cumulative variance over 25 years used in Black's model, giving an estimated average annual volatility of $\sigma = 2.4\%$.[9] Over 25 years this is a relatively modest degree of risk by the usual standards of financial markets, even though it represents a significant risk to pension funds. However, a risk premium of 20 basis points for a volatility of 2.4% per annum is equivalent to a risk premium of 2% for a volatility of 24% per annum. For equities a volatility of 24% is slightly high, though not implausible, whereas a risk-premium of 2% would normally seem low. This suggests that valuing each cashflow at LIBOR-35 would, in this case, underprice slightly the 25-year cashflow.

However, shorter-dated cashflows are likely to be overpriced if valued at LIBOR-35. To see this, consider Figure 5 which shows $Var[\log S(t)]$ for different values of t.[10] The shape of this plot reflects the fact that unanticipated changes in mortality in each year have their effects on $S(t)$ compounded in each subsequent year. If this plot showed, instead, a straight line (as we would find for the Black-Scholes model for equities, which assumes returns follow a random walk) then it would be appropriate to apply the same risk-premium (per annum) to cashflows at all dates. The convex shape in Figure 5 indi-

a buy-and-hold asset rather than as an actively-traded asset.

[7]The functional dependence of $q(t, x)$ on x is a common one used by actuaries for higher ages.

[8]These confidence intervals include allowance for uncertainty in the mean drift of $a(t)$ and $b(t)$.

[9]See, for example, Hull (2003), Section 13.8. The cumulative variance is represented by $\sigma^2 T$ in the extended applications of the Black-Scholes formula proposed by Black. Here $T = 25$ and $\sigma^2 T = 0.014$, implying that $\sigma \approx 0.024$ or 2.4%.

[10]The level of uncertainty reflected in Figure 5 may understate the true level due to model uncertainty. Assessment of the impact of model uncertainty might include consideration of medical rather than statistical issues. As commented above this might account for the differences between the median in Figure 4 and the GAD projection in Figure 2.

cates that short-dated cashflows are relatively low risk and require a smaller risk-premium than long-dated cashflows.

To sum up, we might argue that something close to EIB spot rates might be used to value short-dated cashflows[11] and something over LIBOR-35 for longer-dated cashflows. Therefore, on average, LIBOR-35 represents a reasonable compromise across all cashflows, but it is difficult to judge precisely how good a deal the pension funds got.

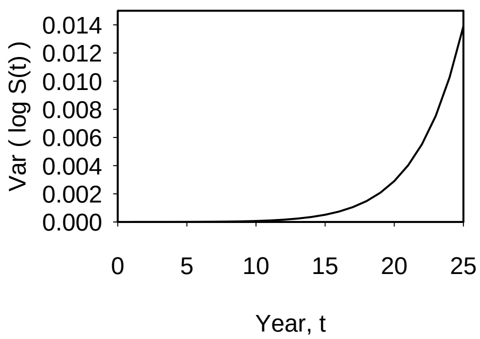

Figure 5: $Var \log S(t)$ corresponding to Figure 4.

Conclusions

The EIB longevity bond represents a pioneering first step to dealing with long-term longevity risk. However, we should be aware of its limitations. It does not provide a perfect hedge against pension plans' individual mortality exposures: there is basis risk between the reference population mortality and the mortality experienced by any individual pension plan. Although we do not have historical mortality tables for pension plans we can investigate graphically the possible degree of basis risk by comparing English and Welsh population mortality with that of UK male assured lives. The improvements in mortality over time for different ages can be seen in Figure 6. It appears from these plots that the trends in mortality improvements over time for England & Wales population mortality more or less matches those for assured lives, except perhaps at high ages. This suggests that basis risk is not as high as one might think, and the interest in this issue indicates that the participating pension funds take this view as well.

[11] Approximately, LIBOR-25 in the secondary market: see Figure 3.

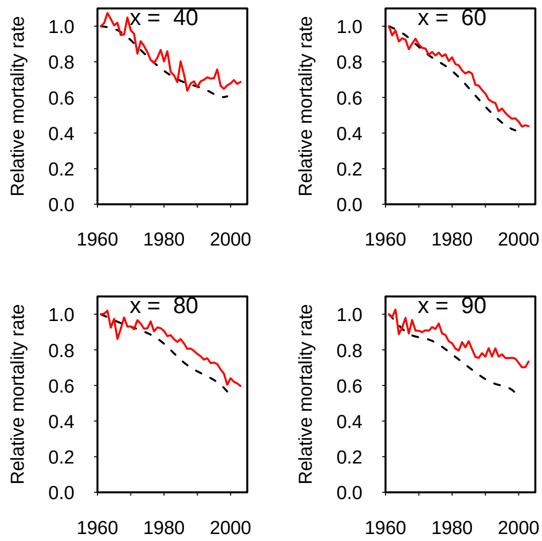

Figure 6: Development of mortality rates over time relative to 1961 values for UK-males assured lives (dashed line) and England & Wales general male population (solid line).

If longevity bonds are to provide effective hedge instruments for the mortality risks actually borne by pension plans then the EIB bond will need to be followed by many others, and these will need to be indexed to the mortality experiences of a much greater range of cohorts.

Acknowledgements

The authors would like to thank Iain Currie and Adrian Gallop for providing the data for Figures 1 and 2 respectively, and Tim Cox, Cliff Speed and David Bowie for their helpful comments on earlier drafts of this article. This article was completed while AC was a visitor at the Isaac Newton Institute, Cambridge.

References

Blake, D., and Burrows, W. (2001) "Survivor bonds: Helping to hedge mortality risk," Journal of Risk and Insurance 68, 339-348.

Cairns, A.J.G., Blake, D., and Dowd, K. (2004) "Pricing death: Frameworks for the valuation and securitization of mortality risk," Working Paper, Heriot-Watt University.

Cox, S.H., Fairchild, J.R., and Pedersen, H.W. (2000) "Economic aspects of securitization of risk," ASTIN Bulletin 30, 157-193.

Currie I.D., Durban, M. and Eilers, P.H.C. (2004) "Smoothing and forecasting mortality rates," Working Paper, Heriot-Watt University.

Hull, J.C. (2003) *Options, Futures, and Other Derivatives (5th Edition).* Prentice Hall.

Lin, Y., and Cox, S.H. (2004) "Securitization of mortality risks in life annuities," Working paper, Georgia State University.