DISCUSSION PAPER PI-0508

Pricing Risk on Longevity Bonds

Andrew Cairns, David Blake, Paul Dawson, Kevin Dowd

October 2005

ISSN 1367-580X

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

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

Pricing risk on longevity bonds

The impact of increasing longevity on pension provision is now a major concern. But better management of longer-term mortality risk is possible following a recent innovative bond issue. We analyse the instrument's risk premium.

ANDREW CAIRNS, DAVID BLAKE, PAUL DAWSON, KEVIN DOWD

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 three-year bond by Swiss Re and Vita Capital that was groundbreaking in this regard, it being 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 the attacks on the World Trade Center in 2001.

Buyers of the Swiss Re 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 than expected in the short term.[1] However, for primary insurers and pension plans, the 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, the most important of which are longer-term mortality risks.

It is now acknowledged that changes in mortality rates over time are only partly 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 past 50 years have been significant. Second, these improvements look random to some extent. Third, the pattern of improvements over time has been different at different ages.

In November 2004, one year on from the issue of the Swiss Re bond, BNP Paribas announced that it had arranged for the European Investment Bank (EIB) to issue a longevity bond 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 is £540 million, and is primarily aimed at UK pension funds. The concept and usefulness of longevity bonds have been discussed for a number of years (Cox et al., 2000, and Blake & Burrows, 2001). But 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), and for both potential issuers and investors to decide that the time is right. As we remark in our conclusions, it is still not clear that the time or the BNP/EIB contract design 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

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'.

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.

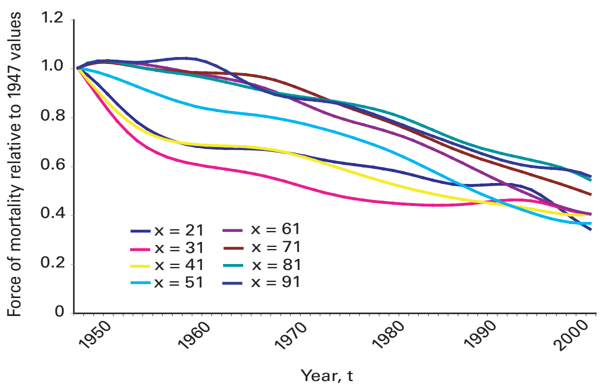

Figure 2. Projected coupons under the EIB longevity bond.

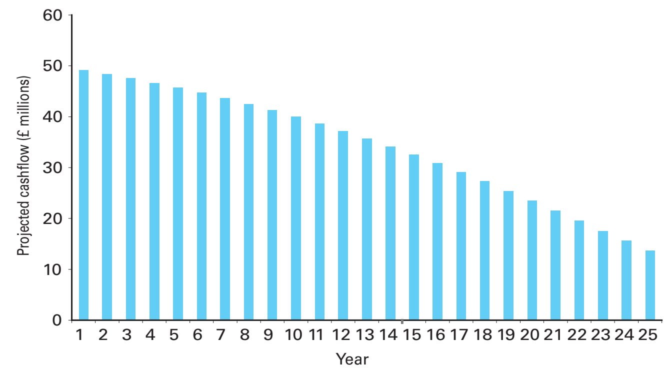

Figure 3. Annualised spot rates on November 18, 2004 for Libor, the EIB (secondary market) and gilt STRIPS.

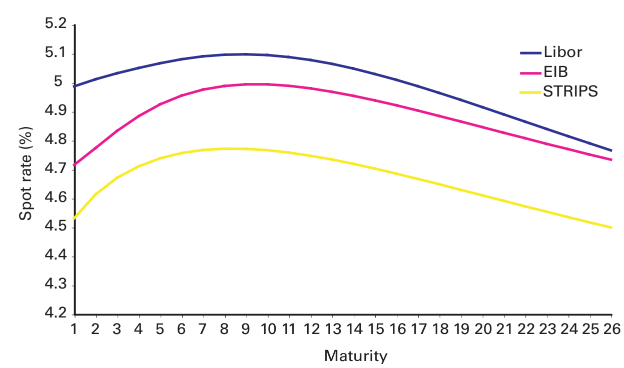

for the index on the issue date. The payment at time t is then $S(t)$ where the values are updated recursively as follows:

1 $$S(t+1)=S(t).(1-m(t-l,65+t))$$

where $m(t,x)$ is the crude (unsmoothed) population central death rate for males aged x in year t for England and Wales. These rates are published regularly 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 offer price for the bond is calculated by taking projected survival rates based on the UK Government Actuary's Department's[4] (GAD) projections. Projected coupons (payable annually) are plotted in figure 2) and these are to be discounted at Libor minus 35 basis points to give the issue price.

Is it a good deal?

The bond is to be issued by the AAA-rated EIB, which will provide excellent security for the coupon payments. The EIB is the guarantor of all payments to the investor. The EIB itself faces some risk associated with differences between experienced mortality and the GAD projection and between projected and actual interest rates because its counterparties, BNP Paribas and Partner Re, are AA-rated. However, this should be of no concern to investors, who face only the EIB.

One might speculate whether the pension plans are being offered 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 (2005). Here, however, we will take a much simpler approach given that we know what the issue price is.

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]

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

3. See www.statistics.gov.uk.
4. See www.gov.uk. The mortality used by the GAD rates (which, by actuarial convention, are usually denoted by q_x or $q(t,x)$ when there is time dependence) are calculated in a different way from the ONS central death rates (by convention, denoted by m_x or $m(t,x)$). The difference, however, is explicitly accounted for in the longevity bond's documentation.
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 for use as a buy-and-hold asset rather than as an actively traded asset.

Figure 4. Median and 95% confidence intervals for future cash flows.

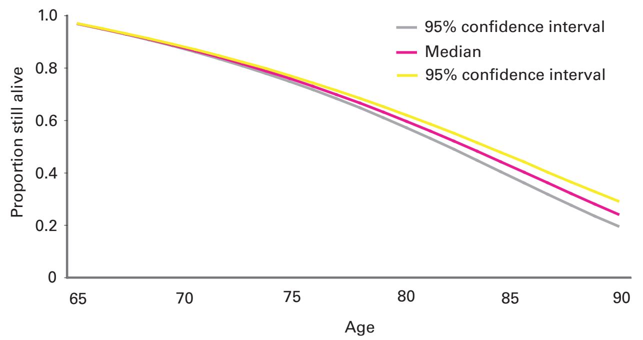

Figure 5. *Var*[log*S*(*t*)] corresponding to figure 4.

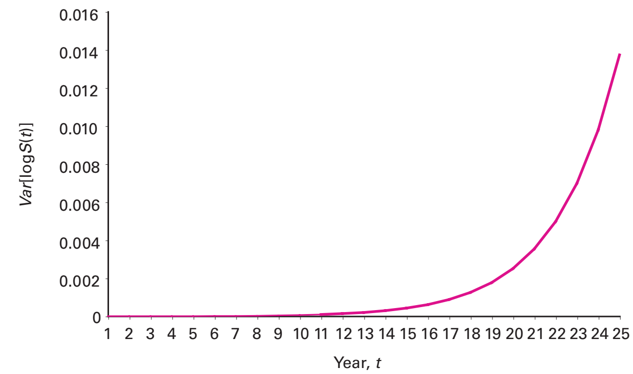

chosen index has a sufficiently large underlying population. A specific aspect of model risk exists in the use of the GAD projection as the benchmark for pricing. If it is believed that this is significantly different from the true mean then the price paid needs to reflect this bias.

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 the death rates $M(t,x)$:

$$m(t,x)=\frac{\exp\big(a(t)+b(t)x\big)}{1+\exp\big(a(t)+b(t)x\big)/2} \tag{2}$$

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. 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 σ=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, slightly underprice 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 a straight line instead (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 indicates 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 and something over Libor-35 for longer-dated cashflows.[11] 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 are being offered.

Conclusions

The EIB longevity bond represents a pioneering first step towards 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 and 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.

However, at the time of writing, despite still being actively marketed, the EIB bond has not generated sufficient demand to be launched. Several reasons have been suggested as to why this might be the case:

- The upfront capital may be too large compared with the risk being hedged, leaving no capital to hedge other risks (such as inflation risk).
- Partner Re is unlikely to be perceived as being a natural holder of UK longevity risk.
- The issue size may be too small to create a liquid market.
- Consultants may be reluctant to recommend it to trustees.
- Fund managers do not currently have a mandate to manage longevity risk.
- Fund managers may not welcome the bond if they believe it would be held passively and they would not make money from it being traded.

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

7. The functional dependence of $m(t,x)$ 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 σ^2T in the extended applications of the Black-Scholes formula proposed by Black. Here T=25 and σ^2T=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.

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 (smoothed) and England and Wales general male population (unsmoothed).

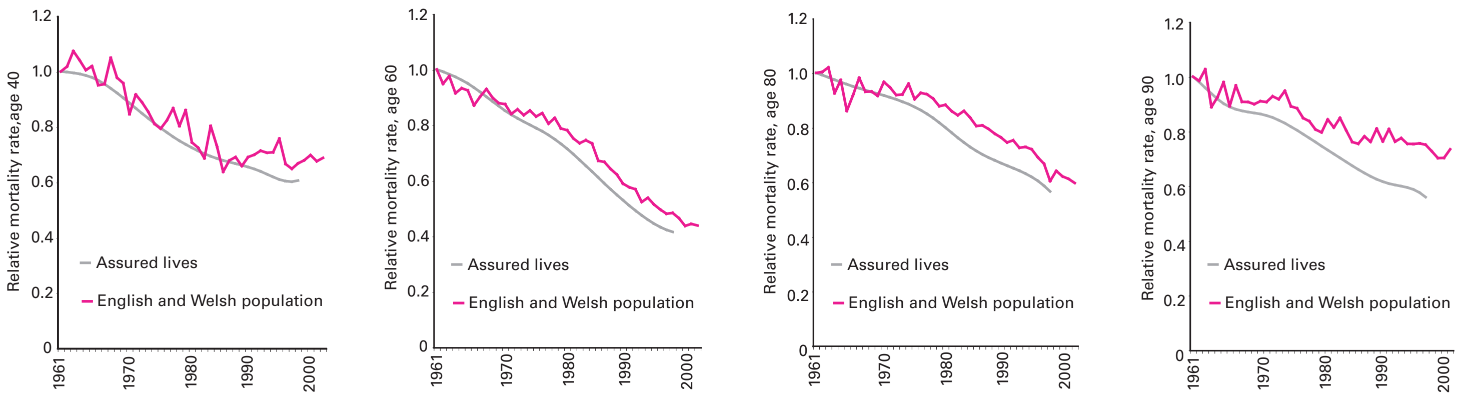

be followed by many others, and these will need to be indexed to the mortality experiences of a much greater range of cohorts. In addition, the problems associated with creating a liquid market in mortality-linked securities need to be resolved. L&P

Andrew Cairns is Professor of Financial Mathematics at the School of Mathematical and Computer Sciences, Heriot-Watt University. David Blake is Professor of Pension Economics at the Pensions Institute, Cass Business School. Paul Dawson is Assistant Professor in the Department of Finance at the College of Business Administration, Kent State University. Kevin Dowd is Professor of Financial Risk Management at the Centre for Risk and Insurance Studies, Nottingham University Business School. The authors would like to thank Iain Currie and Adrian Gallop for providing the data for figures 1 and 2 respectively, and Mark Azzopardi, Tim Cox, Cliff Speed and David Bowie for their helpful comments on earlier drafts of this article. The authors are also grateful to the UK Economic and Social Research Council for financial support under research grant RES-000-23-1036 on Survivor Derivatives. This article was completed while AC was a visitor at the Isaac Newton Institute, Cambridge.

References

Blake D and W Burrows, 2001
Survivor bonds: helping to hedge mortality risk
Journal of Risk and Insurance *68, 339348*

Cairns A, D Blake and K Dowd, 2004
Pricing death: frameworks for the valuation and securitization of mortality risk
Working paper, Heriot-Watt University

Cox S, J Fairchild, and H Pedersen, 2000
Economic aspects of securitization of risk
ASTIN Bulletin *30, 157–193*

Currie I, M Durban and P Eilers, 2004
Smoothing and forecasting mortality rates
Statistical Modelling *4, 279–298*

Hull J, 2003
Options, Futures, and Other Derivatives (5th Edition)
New York: Prentice Hall

Lin Y and S Cox, 2005
Securitization of mortality risks in life annuities
Journal of Risk and Insurance *72, 227252*