

Longevity Bonds: Financial Engineering, Valuation and Hedging

David Blake, Andrew Cairns, Kevin Dowd, and Richard [MacMinn](#)

August 9, 2006

Abstract

This paper examines the main characteristics of longevity bonds (LBs) and shows that they can take a large variety of forms which can vary enormously in their sensitivities to longevity shocks. We examine different ways of financially engineering LBs and consider problems arising from the dearth of ultra-long government bonds and the choice of the reference population index. The paper also looks at valuation issues in an incomplete markets context and finishes with an examination of how LBs can be used as a risk management tool for hedging longevity risks.

1. Introduction

One of the largest sources of risk faced by life companies and pension funds is longevity risk: the risk that members of some reference population might live longer on average than anticipated. For example, if the reference population are annuitants, longevity risk is the risk that annuitants might live longer on average than anticipated in life companies' mortality tables used to price annuities. Longevity risk is an important problem both because of the uncertainty of longevity projections, on the one hand, and because of the large amounts of liabilities exposed to longevity risk, on the other. The uncertainty of longevity projections is illustrated by the fact that life expectancy for men aged 60 is more than 5 years' longer in 2005 than it was anticipated to be in mortality projections made in the 1980s;¹ and the amounts at risk are illustrated by the fact that state and private sector exposure to longevity risk in the UK amounted to £2520bn (or \$4424bn) at the end of 2003 – that is, nearly £40,000 (or \$70,000) for every man, woman and child in the UK.²

Exposure to longevity risk is therefore a serious issue, and yet, traditionally, life companies and pensions funds have had few means of managing it: until recently,

¹ Hardy, 2005, p. 17.

² Pensions Commission, 2005, Fig 5.17, p. 181.

longevity risks were never securitized³ and there were no longevity derivatives that these institutions could use to hedge their longevity risk exposures. However, this state of affairs is changing, and markets for longevity derivatives are starting to develop. Most prominent amongst these are longevity bonds (LBs), which are financial instruments in which payments depend on the realization of a survivor index $S_{t,x}$ for some period t . As its name suggests, the survivor index is the proportion of some initial reference population aged x at time $t=0$ who are still alive at some future time t . If $q_{s,x}$ is the mortality rate between s and $s+1$ for members of the reference population aged x at time $t=0$ and still alive at time s , then the relationship between the sequence $q_{s,x}$ and $S_{t,x}$ is given by:

$$S_{t,x} = (1 - q_{0,x})(1 - q_{1,x}) \dots (1 - q_{t-1,x}) \quad (1)$$

Since we are dealing in this paper with reference populations from a single age cohort, we will simply denote below the survivor index as S_t .

LBs were first proposed by Blake and Burrows, 2001, and the first operational mortality-linked bond (the Swiss Re mortality catastrophe bond) appeared in 2003. A second mortality-linked bond (the EIB/BNP Paribas longevity bond) was announced in 2004 (although it failed to come to market), and various other mortality-linked products have also been issued or are in preparation.⁴

However, many actuaries are still unconvinced that longevity bonds (and related derivatives) will have a significant part to play in the management of longevity risk.⁵ Indeed, even supporters of LBs are divided on some of the key issues. For example, Blake, 2003, and Dowd, 2003, disagree on whether LBs should be issued by the state,⁶ and on the significance of a potential natural excess demand for products that

³ The issues involved in the securitization of longevity risks are discussed further by Cowley and Cummins, 2005 and Krutov, 2006.

⁴ See, e.g., Lane, 2006.

⁵ Evidence of this scepticism was seen in the reaction of many in the audience when a recent paper (Blake *et alia*, 2006) was presented to meetings of the Faculty of Actuaries in Edinburgh and the Institute of Actuaries in London in January and February 2006, respectively.

⁶ Furthermore, governments themselves are reluctant to issue longevity bonds. For example, the UK Debt Management Office, part of the UK Treasury, has recently stated 'the issuance of 'longevity'

are hedges against longevity risk.⁷ Thus, the subject of longevity bonds is both novel and controversial, and much more work remains to be done on it.

Our discussion is organized as follows. In section 2, we discuss the different types of mortality-linked bonds, both extant and hypothetical, and draw out the analogies between LBs (and related securities) and conventional debt instruments. Section 3 discusses how LBs might be financially engineered from existing securities. Section 4 addresses some complications for the basic analytical framework arising from the dearth of ultra-long government bonds and the choice of the reference survivor index. Section 5 considers the valuation of these securities in the presence of market incompleteness. In section 6, we examine the sensitivities of different LBs to longevity (and interest-rate) shocks, and show how these sensitivities allow us to identify specific risk management uses for different types of LB. Section 7 concludes.

bonds was considered but not envisaged. Such instruments would raise broader policy issues – such as the outright transfer of additional longevity risk onto the Government’s balance sheet – that extend beyond a strict interpretation of debt management considerations’ (Debt Management Office, 2004, para 1).

⁷ Blake, 2003, argues that financial institutions, such as pension funds and annuity providers, as a whole are short longevity risk and that there is a shortage of potential private sector issuers with a natural long exposure to longevity risk, such as pharmaceutical companies, owners of long-term care homes, or states and municipalities, such as Florida, which attract retirees. However, a counter-argument made by Dowd, 2003, was that this would result, not in market failure, but in hedges against longevity risk selling at a premium. Furthermore, as longevity products become securitized, it is arguable the way will become open for capital markets institutions to take on longevity risk themselves, and longevity risk exposure is attractive to such institutions because of its low beta with respect to more conventional financial risk factors. Brown and Orszag, 2006, also question whether governments should issue longevity bonds given their existing extensive exposure to longevity risks via public pension systems, although they recognise that the state does have a potential role in the intergenerational sharing of longevity risks. The UK Pensions Commission, 2005, p45, argues that there is a case for government issuance, not in net terms, but only if the government removes itself from some of the exposures it is currently assuming. In particular, it should stop insuring the longevity risk of the young working generation by having a state pension age that is fixed over long periods of time and hence independent of increases in longevity. The young are natural hedgers of this risk through their ability to extend their working lives, the Pensions Commission argues.

2. Types of mortality-linked bond⁸

We can examine existing or proposed bonds to get some idea of the characteristics of longevity bonds that issuers and investors might find attractive.

2.1 The Swiss Re mortality catastrophe bond⁹

The first bond with cash flows linked to the realisation of a composite mortality index, M_t , was the Swiss Re bond issued in December 2003. This bond had a maturity of three years, a principal of \$400m, and offered investors a floating coupon of LIBOR+135 basis points. In return for this coupon rate, the principal repayment is dependent on the realized value of a weighted index of mortality rates in five countries, M_t . The principal is repayable in full only if the mortality index does not exceed 1.3 times the 2002 base level during any year of the bond's life, and is otherwise dependent on the realized values of the mortality index.¹⁰ The bond was issued through a special purpose vehicle (SPV) called Vita Capital: this was convenient from Swiss Re's point of view because it kept the cashflows off-balance sheet, but also helped to reduce the credit risk faced by investors.¹¹

The main characteristics of this bond can be summarized as follows:

- The bond was designed to be a hedge to the *issuer*.
- The issuer *gains* if M_t is extremely *high* (and conversely, the buyer gains if M_t is not extremely *high*).

⁸ We adopt the term mortality-linked bond to represent the general class of bond whose cash flows are linked to realised mortality, and we differentiate between mortality bonds, whose cash flows, $f_t(M_t)$, are linked to a mortality index, M_t , and longevity bonds, whose cash flows, $f_t(S_t)$, are linked to a survivor index.

⁹ The first bond that we consider is not, in fact, a longevity bond intended to hedge longevity risk. Rather it is a mortality catastrophe bond designed to hedge brevity risk. Brevity risk is the risk of too short a life and so from the insurer's perspective it is the risk that benefits on life policies will have to be paid before expected. Nevertheless, an understanding of the bond's characteristics will provide a useful guide to the potential design of longevity bonds. Brevity risk is discussed in more detail in MacMinn and Richter (2004).

¹⁰ That is, the bond is a principal-at-risk bond.

¹¹ More details on this bond and the later EIB/BNP Paribas bond are given in Blake *et alia*, 2006. There was another issue of the Swiss Re bond via an SPV called Vita Capital II in April 2005. For details on this latter issue see <http://www.artemis.bm/html/dealdir/index.htm>.

- The first two points together imply that the bond is a hedge against a portfolio dominated by life insurance/reinsurance (rather than annuity) policies.
- The bond is a *short-term* bond designed to protect the *issuer* against an *extreme increase* in mortality, such as that associated with an influenza pandemic.
- The mortality index, M_t , is a weighted average of mortality rates over 5 countries, males and females, and a range of ages.
- The bond is a standard coupon-plus-principal bond in which the coupons float with LIBOR and only the principal is at risk from a mortality deterioration that might occur during the period until the bond matures, and it is the spread over LIBOR that compensates the holder for allowing the principal to be at risk.

The precise payment schedules are given by the following $f_t(\cdot)$ functions:

$$f_t(\cdot) = \begin{cases} LIBOR + spread & t = 1, \dots, T-1 \\ LIBOR + spread + \max(0, 100\% - \sum_i L_i) & t = T \end{cases}$$

(2)

where L_t is the following loss function:

$$L_t = \begin{cases} 0\% & M_t < 1.3M_0 \\ [(M_t - 1.3M_0)/(0.2M_0)] \times 100\% & \text{if } \begin{cases} 1.3I_0 \leq M_t \leq 1.5M_0 \\ 1.5I_0 < M_t \end{cases} \text{ for all } t \\ 100\% & \end{cases}$$

and where M_0 is the base mortality index and M_t is the mortality index for year t . The coupon payments involve no dependence on any mortality index, whilst the principal repayment is a piecewise linear function of the mortality index.

2.2 The EIB/BNP Paribas longevity bond

The second bond was much closer in nature to the ‘classical’ survivor bond proposed by Blake and Burrows, 2001. This bond was announced by the European Investment Bank (EIB) in November 2004. It had an initial value of £540m, an initial coupon of £50m and a maturity of 25 years. The structurer/manager was BNP Paribas. The longevity risk was to be reinsured through the Bermuda-based reinsurer Partner Re

which contracted to make annual floating rate payments (equal to $\pounds 50m \times S_t$) to the EIB based on the realized mortality experience of the population of English and Welsh males aged 65 in 2003 (published by the UK Office for National Statistics) and receive from the EIB annual fixed payments based on a set of mortality forecasts for this cohort. The mortality forecasts were based on the UK Government Actuary's Department's 2002-based central projections of mortality, adjusted for Partner Re's own internal revisions to these forecasts. Since the EIB also wished to pay floating rate in euros, this arrangement was then supplemented by a cross-currency (i.e., fixed-sterling-for-floating-euro) interest-rate swap between the EIB and BNP Paribas.¹²

The main characteristics of this bond are therefore:

- The bond was designed to be a hedge to the *holder*.
- The issuer *gains* if S_t is *lower* than anticipated (and conversely, the buyer gains if S_t is higher than anticipated).
- Thus, the bond is a hedge against a portfolio dominated by annuity (rather than life insurance/reinsurance) policies.
- The bond is a *long-term* bond designed to protect the *holder* against *any* unanticipated *improvement* in mortality up to the maturity date of the bond.
- S_t involves a single national survivor index.
- The bond is an annuity (or amortizing) bond and all coupon payments are at risk from longevity shocks. More precisely, the payment schedules are directly proportional to the survivor indices:

¹² It is also worth noting that there are a number of actual or potential credit exposures in this arrangement. To begin with, the end investors are exposed to the risk of default by the LB issuer, the EIB, and the EIB is backed by a AAA credit rating. The EIB has a commitment to make longevity-linked payments in sterling, and engages in a swap with BNP to convert this commitment for one to make floating euro payments. This swap means that the EIB and BNP are then potentially exposed to each other. At the same time, BNP takes on longevity exposure which it hedges with Partner Re, so BNP takes on a credit exposure to Partner Re. In terms of protection, the EIB has the protection of BNP's commitment to take on the bond's longevity exposure, and this commitment is backed by BNP's AA credit rating and the knowledge that BNP has reinsured this risk with Partner Re. And, for its part, BNP has the protection of the EIB's AAA credit rating, on the one hand, and the reinsurance provided by Partner Re, whose rating is AA.

$$f_t(S_t) = £50m \times S_t \text{ for } t=1,2,\dots,T; T=25 \quad (3)$$

2.3 Characteristics of longevity bonds

In the light of the above analysis, any design of longevity bonds must take into account the following characteristics:

- Whether the bond is issued or held as a hedge.
- The type of portfolio that is hedged by the bond, i.e., whether the portfolio is predominantly life assurance contracts or predominantly annuities.
- The type of bond: coupon-plus-principal, annuity, etc.
- The survivor index used.
- The nature of the payment function and the way in which it is contingent on S_t , i.e., what form does $f_t(S_t)$ take?

It is clear that LBs can vary across many different dimensions (e.g., type of bond, institution and position to be hedged, maturity, survivor index, credit risks involved, specification of $f_t(S_t)$, etc.).

We can also envisage many other types of longevity bond with different characteristics:

- Longevity zeros (LZs): these are the equivalents of conventional zeros, and have similar purposes, e.g., as building blocks for more complicated tailor-made securities
- Survivor bonds: these continue to make payments for as long as any member of the reference population is still alive. These LBs have a stochastic maturity equal to the time of death of the last survivor from the reference population. The attraction of this open-endedness is that these bonds can provide a better hedge to an annuity book than a LB that matures whilst members of the reference population are still alive.¹³
- Principal-at-risk LBs: these are LBs whose coupons might be fixed or interest-sensitive, but whose principal repayments are functions of a survivor index.

¹³ This was the form of the bond originally suggested by Blake and Burrows, 2001.

- Inverse LBs, in which the $f_t(S_t)$ are inverse functions of S_t . A simple example might be the payment function $f_t(S_t) = k(1 - S_t)$ for some $k > 0$. These are comparable to conventional inverse floaters whose coupon payments move inversely with market interest rates, and can be regarded as a form of longevity structured note. Unlike standard LBs whose coupon payments fall over time, the coupon payments on inverse LBs rise steadily over time.
- Collateralized longevity obligations (CLOs) comparable to conventional collateralized debt obligations (CDOs). In the same way that a CDO is a tranche of a pool of debt instruments, a CLO would be a tranche on a pool of LBs. Different tranches would have different exposures to longevity risk (e.g., the first tranche might absorb the first 5% of any gains or losses, the second tranche might absorb the second 5%, and so on). Different tranches would therefore have very different risk exposures, and correspondingly different expected returns.¹⁴ The CLO could also simply be linked to a portfolio of life annuities.

3. Financial engineering longevity bonds

There are four ways in which an investment bank might consider constructing LBs. The simplest way is by decomposing the cash flows on a conventional bond. A more complex way involves the combination of LZs and a longevity swap. A third way combines LZs with a series of forward contracts. A fourth way involves a conventional long term bond with options on the toxic tail risk. All methods have advantages and disadvantages.

¹⁴ We can also envisage synthetic CLOs that correspond to well-known synthetic CDOs, which are tranches on pools of credit-default swaps. A synthetic CLO would therefore be a tranche on a pool of longevity swaps. These synthetic derivatives have the attraction of requiring less upfront capital, and this enables holders to increase their leverage.

3.1 Engineering a longevity bond by decomposition

The first approach works as follows. Consider a government annuity bond¹⁵ paying a fixed annual coupon of unity and having a fixed maturity of T which is greater than the maximum expected life of the population group underlying the survivor index. A bank could put this bond into an SPV and then ‘split’ claims on the SPV into two survivor-dependent instruments, a LB that pays a coupon equal to S_t in year t for $t = 1, \dots, T$, and an inverse LB (ILB) that pays a coupon equal to $1 - S_t$ in year t , as shown in Fig. 1.

Insert Figure 1 here

The principal advantage of this method of financially engineering LBs is that there is no credit risk involved. This is because there will always be sufficient funds in the SPV to continue making payments on the LB, even in the extremely implausible case of no deaths prior to T . The value of the ILB therefore represents an upper bound to the cost of the absolute guarantee that the LB will make all payments in full. However, the downside of this method is that it produces a by-product (in this case, the ILB) that may be problematic (e.g., the ILB might be relatively unattractive to investors compared with the LB).

3.2 Engineering a longevity bond using longevity zeros and a longevity swap

A second approach is to use a series of T zero-coupon bonds (LZs) of increasing maturity combined with a longevity swap (sometimes also known as a mortality or survivor swap).¹⁶ A longevity swap (LS) is a swap involving an exchange of one or more payments over a set term, at least one leg of which is linked to the realized value of a survivor index (Cox and Lin, 2004, Lin and Cox, 2005, or Dowd *et al*, 2006b). The bank then sets up an SPV consisting of the LZs and the LS, and this SPV provides the desired longevity bond.

¹⁵ Where not directly available, such a bond can easily be constructed from a standard T -year bullet bond if a strips market in the bullet operates.

¹⁶ This was how the EIB/BNP Paribas bond was actually constructed. For more details on its construction, see Blake *et alia*, 2006, section 4.3.

The floating leg of the swap involves a payment S_t in year t , contingent on the realized survivor index, and this is the amount to be paid to holders of the LB. The other leg of the swap involves a fixed payment that is set at time 0; this amount could, for example, be based on \hat{S}_t , the (real-world) expected value calculated at time 0 of the survivor index in year t . To ensure that the two payment legs have the same initial value (and hence ensure that the swap itself has a zero starting value), the fixed payments would incorporate a premium π that might be positive, negative or zero: the fixed payment would therefore be $(1 + \pi)\hat{S}_t$, rather than \hat{S}_t . The valuation of π is considered in section 5 below. The net payoff to the SPV on the LS is $S_t - (1 + \pi)\hat{S}_t$ in year t . The LZ maturing in year t pays $(1 + \pi)\hat{S}_t$. The combination of the LZs and the LS has exactly the same payment schedule as the desired LB, namely S_t in year t :

$$(1 + \pi)\hat{S}_t + (S_t - (1 + \pi)\hat{S}_t) = S_t \quad (4)$$

The cashflows from this arrangement are illustrated in Fig. 2.

Insert Figure 2 here

Other types of LB can be constructed in similar ways. For example, LZs can be constructed by utilizing the same sort of SPV we have just discussed, but instead of selling the claims to the different payments as a package, the bank sells them off individually instead. A T -year, fixed-interest principal-at-risk LB can be constructed by combining a series of T ZCs (which accounts for the coupon payments) with a T -year maturity LZ with a higher face value, to account for the at-risk principal payment. For its part, a T -year inverse LB could be constructed using an SPV with a series of T ZCs, with face values of $(1 - (1 + \pi)\hat{S}_t)$ for each t and a LS paying a floating leg of S_t and receiving a fixed leg of $(1 + \pi)\hat{S}_t$. And, finally, a CLO can be constructed using an SPV whose assets consist of a collection of LBs, the claims on which are sold off in the form of tranches with different priorities on any profits or losses made by the SPV.

The key advantage of this method of creating a LB is that there is no (possibly undesired) by-product, such as an ILB. The disadvantage is the credit risk associated with the LS.¹⁷

3.3 Engineering a longevity bond using longevity zeros and forward contracts

A third way is again to use a series of T zero-coupon bonds (LZs) of increasing maturity (with the bond maturing in year t paying $(1 + \pi_t)\hat{S}_t$, with \hat{S}_t as defined above), but this time the LZs are combined with a series of forward contracts to exchange the survivor payoff S_t for the forward price $(1 + \pi_t)\hat{S}_t$ in year t . The net payoff on the forward contract is $S_t - (1 + \pi_t)\hat{S}_t$ in year t , so the combined payoff from the LZ and forward contract maturing in year t is:

$$(1 + \pi_t)\hat{S}_t + (S_t - (1 + \pi_t)\hat{S}_t) = S_t \quad (5)$$

The risk premium in the forward contract, π_t , is set at a level that ensures that the value of the forward for each t has zero value at time 0. The cash flows for the SPV established for the purpose will be similar to those in Figure 2.

The method proposed here would be equivalent to the approach involving swaps if all the risk premia were the same, but the important point to be made is that these premia need not be the same. One of the advantages of the third approach is that different pieces of the risk can be held by different market participants. Further, if the forward contracts were replaced by exchange traded futures contracts then any credit risk problems would be alleviated even more. Of course, to the extent that the risks are held in forwards by different participants the credit risk might also be alleviated.

3.4 Principal-at-risk longevity bonds

¹⁷ This is however a relatively minor disadvantage, and the parties involved can handle it using standard credit enhancement methods. Because credit enhancement is now well understood by practitioners, we will ignore credit issues in the remainder of our discussion.

It is possible that an annuity provider only wishes to hedge the tail longevity risk. It is the tail that poses the greatest risk of projecting longevity incorrectly. Turner, 2006, notes that errors in forecasts have been significant. For UK data and the Government Actuary's Department (GAD) 1983 forecast, he noted that the actual 1994 mortality rate of a 65 year old male was 17% below the 1983 projected rate; he also noted that, using the 2003 GAD forecast, the 2014 projected mortality rate for a 65 year old male was 46% below the 1983 forecast.

If a 10-15 year period yields reasonably reliable mortality rate predictions, as Turner suggests, then any bond issue with a maturity date beyond 15 years might involve excessive risks from the issuer's perspective. Suppose, however, that an insurer issues a 20-year bond with a sequence of floating or fixed coupons and that the principal repayment is put at risk in a manner similar to the Swiss Re design for Vita Capital. The idea would be to provide some payoff or sequence of payoffs if longevity were badly underestimated. It should not matter whether the contingent payoff is a single payment or a sequence; it would be the value of that loss payoff that would matter. The construction of the loss would differ from Vita Capital because the trigger(s) would utilize a survivor index but otherwise could be adjusted to be any increasing function of the survivor index designed to cover some of the difference between of actual loss and the expected loss in each of the last five years.

Suppose the bond is a standard coupon-plus-principal bond in which the coupons float with LIBOR and, as in the Vita Capital case, only the principal P is at risk, in this case from longevity improvements that might occur during the period until the bond matures, and it is the spread over LIBOR that compensates the holder for allowing the principal to be at risk. Let $L_t(S_t)$ denote the loss in period t due to an index sufficiently large that it triggers a reduction in principal. The payment schedules are given by the following:

$$f_t(\cdot) = \begin{cases} LIBOR + spread & t = 1, \dots, T-1 \\ LIBOR + spread + \max \left\{ 0, \left(P - \sum_{t=16}^{20} L_t(S_t) \right) \right\} & t = T \end{cases} \quad (6)$$

Another example might be a 25-year bond with a (real-world) expected value for the survivor index in year 25 of, say, \hat{S}_{25} . If S_{25} turns out to be higher than this, the principal is reduced in proportion to the factor \hat{S}_{25}/S_{25} . This provides some protection for the issuer of the bond, say an annuity provider, against the so-called ‘toxic tail’: that is, the risk associated with a larger proportion of policyholders surviving well beyond 25 years than anticipated.

4. Complications

We turn now to address certain real-world complications to the basic framework outlined above.

4.1 Availability of government bonds of sufficient maturity

The ability to construct LBs is constrained by the availability of suitable conventional instruments on which the financial engineers can go to work, and this means in practice that the maximum maturity of longevity bonds is limited to that of available government debt.¹⁸ This can be a problem in the case of countries where a life company or pension fund might be concerned about payments at horizons well beyond those of existing government debt, and might therefore want to buy a LB whose maturity exceeded that of such debt.¹⁹

¹⁸ In principle, we could also financially engineer longevity-dependent bonds from non-government debt, but this would require appropriate credit-enhancement mechanisms to manage the extra credit risks involved. Nor does it really solve the problem discussed in the text, because the maximum maturity of private sector debt is not as long as that of government debt.

¹⁹ That this is a significant practical and not just a hypothetical problem is suggested by the empirical work of Cairns *et alia*, 2006b, and Dowd *et alia*, 2006a: their work indicates that a considerable amount of the risk in annuity books (which has been described as the ‘toxic tail’) arises from very long-term horizons (i.e., well beyond 25 years). So, for example, a LB with a horizon of 40 or 50 years will provide a considerably better hedge for an annuity book based on a 65-year old reference population than a LB predicated on a 25-year horizon. LBs with ultra-long horizons are therefore much better hedges than LBs limited to only 25 years.

This line of reasoning translates into an argument for the relevant governments to issue ultra-long debt. Fortunately, governments are beginning to realize the funding attractions of ultra-long debt, and the first ultra-long government debt in modern times – a 4% bond with a face value of €6bn with a maturity of 50 years – was issued by the French Government in February 2005. The UK Government also issued two 50 year bonds in 2005 (in July and December) each with a nominal value of £2.25bn and a coupon of 4.25%. It also became the first government in the world to issue 50-year inflation index-linked (IL) bonds: in September and October 2005, it issued £1.25bn and £0.675bn, respectively, of 1.25% IL bonds. Provided that governments issue enough such debt, we can foresee that financial institutions will be able to use that debt to construct synthetic longevity bonds with maturities out to 50 years, and given likely life expectancies, a 50-year horizon is long enough to eliminate almost all an annuity provider's exposure to longevity risk.

4.2 Survivor index problems

The choice of survivor index is critical to the success of longevity bonds. The bond's cash flows must provide a reasonably close match for the payments the hedger needs to make if the bond is to provide an effective hedge. This suggests that a large number of longevity bonds might need to be issued with survivor indices covering the full age and gender spectrum. On the other hand, a viable market in longevity bonds needs to attract sufficient speculative demand. Speculators create liquidity but a liquid market in longevity bonds is likely only if a small number of longevity bonds are traded. The choice of survivor indices therefore needs to balance hedging needs against speculative interest.

Survivor indices also suffer from other problems (discussed in more detail in Blake *et alia*, 2006; also see Stallard 2006) that need to be addressed if market participants are to have confidence in longevity bond issues:

- Survivor indices are constructed from mortality data that is published infrequently and subject to incurred-but-not-reported (IBNR) errors. For example, there can occasionally be a significant delay between death

occurrence and death registration. Deaths of British citizens occurring abroad are never counted in UK mortality statistics.

- The historical data going into the index typically needs to be smoothed. The methods used to smooth or graduate the crude mortality data change from time to time; changes in the calculation methodology introduce uncertainty (or a possible perception of uncertainty), which may put off potential investors.
- They are subject to integrity and contamination risk. The index underlying a mortality-linked security needs to have, and must be perceived to have, integrity in the way that it is calculated. i.e., it must be based on accurate and complete mortality data. For example, age at death can be unreliably reported, particularly at higher ages. Deaths of visitors to the UK are counted in UK mortality statistics and these will contaminate an index designed to hedge a UK pension fund's exposure to longevity risk.
- There are issues of moral hazard. Moral hazard can arise, for example, when data providers have much earlier access to the data than investors. This type of problem might not affect the attitude of long-term investors (although it might affect the price they are prepared to pay) but it is likely to put off short-term investors. Moral hazard can also exist when there is the possibility that the underlying index might be manipulated (e.g., it can arise where firms have an incentive to misreport mortality statistics to influence the value of the survivor indices that determine the payments on longevity bonds).
- They involve projections of future mortality that are subject to both model and parameter risk and an accepted and transparent mortality forecasting model needs to be used to determine the 'official' values of survivor indices.

5. Valuation

Longevity securities involve significant valuation problems. Whereas conventional fixed-income securities can be valued using the standard spot yield curve and zero-

arbitrage (or net present value) methods, this is not possible with longevity securities because of market incompleteness. We consider two approaches to this problem.²⁰

5.1 Distortion approaches to pricing

One solution to this problem is to use a distortion approach such as the Wang Transform (Wang, 1996, 2000, 2002, 2003). This approach distorts the distribution of the survivor index to create risk-adjusted expected values (or certainty equivalents) that can be discounted at the risk-free rate. The extent of the risk adjustment should reflect the market prices of risk for other assets in the market place which permit trading of other incomplete market risks. More specifically, if $\Phi(\cdot)$ is the standard normal cumulative density function, the Wang distortion operator is:

$$g_\lambda(u) = \Phi[\Phi^{-1}(u) - \lambda] \quad (7)$$

where $0 < u < 1$ is a cumulative probability and the parameter λ is the market price of risk. If an instrument produces a random cash flow Y with distribution function $G(y)$, then

$$G^*(y) = g_\lambda(G(y)) \quad (8)$$

can be interpreted as its ‘risk-adjusted’ cumulative distribution function, and the ‘fair’ value of the instrument is the mean of Y under $G^*(y)$ discounted at the risk-free rate.

²⁰ A third approach, not illustrated here, adapts financial economic theory for incomplete financial markets (Froot and Stein, 1998). They discuss how the capital structure of a firm, and the nature and size of a specific transaction affects the price for the risk transferred.

Let $E^*(\cdot)$ be the expectation operator associated with the Wang transformed distribution $G^*(\cdot)$ in (8), then the value of the LB, $V(LB)$, is given by:

$$V(LB) = \sum_{t=1}^T D_0^t E^*(S_t) \quad (9)$$

where D_0^t is the risk-free discount factor at time 0 for fixed payments at time t (that is, the price at time 0 for \$1 payable with certainty at time t) and $E^*(S_t)$ is the expected cash flow. The value of an ILB ($V(ILB)$) can then be determined using the following zero-arbitrage relationship linking the values of an annuity bond (AB), a LB and an ILB:

$$V(AB) = V(LB) + V(ILB) \quad (10)$$

where $V(LB)$ can be determined using standard present value methods by discounting the coupons using the appropriate spot yields.

The risk premium π in a longevity swap can also be determined using the Wang Transform. Suppose the swap value is given by:

$$\text{Swap value} = V(S) - V((1 + \pi)\hat{S}) = V(S) - (1 + \pi)V(\hat{S}) \quad (11)$$

where $V(S)$ is the value of the floating leg payments received and $V((1 + \pi)\hat{S})$ is the value of the fixed leg payments paid. We now calculate the value of fixed payments $V(\hat{S})$ using standard present-value methods, and calculate the value of the floating payments $V(S)$ as the expected value of the S_t cash flows under the Wang transformed probability measure (7) discounted at the risk-free rate. The premium is then set so that the swap value is zero, viz.:

$$\pi = \frac{V(S)}{V(\hat{S})} - 1 \quad (12)$$

Applying the expectations operator under the Wang Transform to (5) yields the following value for the LZ and forward contract method of engineering a LB:

$$\begin{aligned} V(LB) &= \sum_{t=1}^T D_0^t E^* \left(\hat{S}_t + (S_t - \hat{S}_t) \right) \\ &= \sum_{t=1}^T D_0^t E^* (S_t) \end{aligned} \quad (13)$$

The risk premium in the case of the approach using forward contracts does not require a constant risk premium. There could be a different risk premium, $\pi_t = E^*(S_t)/\hat{S}_t - 1$, for each forward contract because while the SPV would be on one side of all the forwards, different counterparties could be involved with each contract on the other side. However, in the absence of credit risk, the price of the LB under all three financial engineering methods will be identical.

5.2 Risk-neutral pricing

An alternative is to use the risk-neutral approach to pricing favoured by many recent authors (see, for example, Milevsky and Promislow, 2001, Dahl, 2004, Dahl and Møller, 2005, Miltersen and Persson, 2005, Cairns *et alia*, 2006a – and references therein – and Cairns *et alia*, 2006b). This is based on long-established financial economic theory that states that even in an incomplete market, if the overall market is arbitrage free, then there exists at least one such risk-neutral measure Q that we can use to calculate fair prices. The problem is then to identify what this might be, given the paucity of reference securities against which we can calibrate a risk-neutral pricing measure.

An example of this type of approach is suggested by Cairns *et alia*, 2006b (which is explained in more detail in Appendix A and is used to produce the results discussed below in section 6). They assumed that the market price of longevity risk is constant

and estimated it using the market price of longevity risk proposed in November 2004 for the EIB/BNP longevity bond. This approach is simple to implement and one can argue that more sophisticated assumptions about the dynamics of the market price of risk are pointless given that, at the time, there was just a single item of price data available for a single date (and even that is no longer valid²¹). However, as a market in longevity-dependent securities becomes established over time, we will begin to collect hard pricing data against which we can calibrate our models and evaluate the adequacy of alternative assumptions about the market price of risk. In the meantime, the choice of Q is a matter of speculation that is best guided by a combination of sound economic judgement and straightforward assumptions that are consistent with whatever data are available.

If we now assume that the risk-neutral measure Q has been established then a standard LB that pays S_t at time t for $t = 1, \dots, T$ has a price at time 0 of:

$$V(LB) = \sum_{t=1}^T D_0^t E_Q(S_t / \Omega_0) \quad (13)$$

where Ω_0 represents the information about mortality rates that is available at time 0 and $E_Q(S_t / \Omega_0)$ is the expected value of S_t under the risk-neutral measure Q , conditional on Ω_0 . The risk-neutral approach to pricing can also be applied to the LB engineered using swaps (equation (4)) and forwards (equation (5)).

6. Sensitivities and Hedging Uses of Longevity Bonds

It is interesting to investigate the sensitivities of some of these different LBs in the face of shocks to the mortality rate (q) and interest rates (r). These sensitivities give a

²¹ The reasons for the bond being withdrawn appear to do mainly with design rather than pricing issues. The principal problem was that it was too generic and did not adequately hedge the basis risk in pension fund and annuity book liabilities (which contain both male and female members of different ages, not drawn from the national population, but from the select group of generally longer living people who join pension schemes and purchase annuities).

first-order approximation to their hedging features,²² and can be measured using the elasticities of LB values with respect to changes in each of these two rates.²³ The elasticity of the LB values with respect to S_t can be inferred from (1).

To illustrate some of the elasticity possibilities, we consider the following four LBs:^{24,25}

- Longevity zero where the payment is equal to S_T .
- Standard LBs where the payments are equal to S_t for each t from 1 to T .
- Principal-at-risk LBs where the principal is equal to S_T , with coupons adjusting over time in line with market interest rates.²⁶
- Principal-at-risk LBs where the principal is equal to S_T , with fixed coupons set equal to the current market interest rate at the time of issue.

For each of these, we also consider horizons T varying from 1 to 50.²⁷

²² One should keep in mind that these first-order approximations to hedging might be quite inaccurate, and also ignore complications like basis risk due to differences between the reference population underlying the LB and the population underlying the position to be hedged. However, we are trying to form a broad sense of the hedging features of these securities, and are not trying to work out ‘optimal’ hedging strategies.

²³ They can also be measured in other ways too: for example, the sensitivities of LBs to interest rate changes can be measured using duration or PV01 methods. However, we prefer to use elasticities here because they can easily be used to measure sensitivities to q as well as to r , and elasticity measures make for more straightforward comparisons of sensitivities because they are independent of the value of the position.

²⁴ We do not consider more sophisticated securities such as CLOs and synthetic CLOs, but the 4 LBs considered here provide a good illustration of the very wide range of elasticity patterns possible with LBs. Examination of the properties of CLOs and synthetic CLOs would be a good topic for a later study.

²⁵ The elasticities of the inverse equivalents to each of these, of the form $1 - S_t$, are discussed in Appendix B.

²⁶ We do not consider principal-at-risk LBs of the Swiss Re type, in which the principal is at risk from an extreme mortality event. This is because LBs are tailored to specific concerns about extreme mortality, and will tend to be less sensitive to mortality shocks than the principal-at-risk considered here precisely because they only ‘kick in’ when extremely low-probability events occur.

²⁷ We also restrict ourselves to considering long positions in these securities. Results for short positions can be obtained by multiplying the signs of all the long-position elasticities by -1.

We use the calibrated mortality forecasting model of Cairns *et alia*, 2006b, to price the security.²⁸ We then shock q and r in turn, and for each T from 1 to 50 estimate the q and r elasticities:

$$\eta_{q,T} = \left. \frac{\Delta V / V}{\Delta q / q} \right|_T \quad (q \text{ elasticity for period } T) \quad (14a)$$

$$\eta_{r,T} = \left. \frac{\Delta V / V}{\Delta r / r} \right|_T \quad (r \text{ elasticity for period } T) \quad (14b)$$

where V is the price of the relevant security.²⁹ In the examples below, all of the elasticities have been calculated numerically. However, in the case of LZs (Section 6.1), simple expressions can be written down for the q and r elasticities.

6.1 Elasticities for longevity zeros

We begin with LZs. Recall that the risk-neutral price of a period T LZ is given by $E_Q(S_{T,x})$. It is easy to show that their q and r elasticities are given by:

$$\eta_{q,T} = -E_Q \left(\left(\sum_{s=0}^{T-1} q_{s,x} \right) S_{T,x} \right) / E_Q(S_{T,x}) \quad (q \text{ elasticity for period } T) \quad (15a)$$

$$\eta_{r,T} = -rT \quad (r \text{ elasticity for period } T) \quad (15b)$$

where $q_{s,x}$ is mortality rate for an individual aged $x+s$ at time s . Eqn (15) indicates that the elasticities have ‘nice’ intuitive properties. Plots of these elasticities against T are given in Figure 3. The longevity elasticity starts at 0 and falls as T gets larger, and

²⁸ The model is calibrated for English and Welsh males aged 65 in 2003 using data from the UK Government Actuary’s Department. More details about the model and its calibration are provided in Appendix A.

²⁹ To calculate these elasticities we first calculate the initial price V ; we then shock both q and r upwards by 1% (i.e., we set $\Delta q / q$ and $\Delta r / r$ equal to 1%), estimate the resulting price change and use (14a) and (14b) to obtain the elasticities. The interest rate is shocked from the assumed value of 0.05, and for convenience the spot rate term structure is assumed to be flat. Mortality rates are shocked away from the values generated by the mortality forecasting model. More specifically a 1% increase in mortality rates means that each of the future random $q_{t,x}$ are replaced by $1.01q_{t,x}$.

then falls rapidly when T gets high: this reflects the cumulative effect of the force of mortality as T gets larger, as given in (15a). For its part, the r elasticity also starts at 0, but falls more gradually (and linearly) as T gets larger, in line with (15b). In other words the asset loses value if either of q or r rise.³⁰ In a wider context, we might consider the q and r elasticities on the liabilities side (taken at fair value) of the balance sheet. A book of annuities is, in effect, a collection of LZ's so the liabilities will fall if either of q or r rise. In contrast, the fair value of a book of life insurance business will fall if q falls or r rises.

Insert Figure 3 here

In terms of absolute values, LZs (and their inverse equivalents) typically have the largest elasticities of any of the securities considered here, and this means that these securities (typically) provide greater leverage than the other securities when used as hedge instruments.

In addition, since the LZ has a notable interest rate exposure, any hedging strategy must also take account of its impact on the user's interest-rate risk exposure. However, since interest-rate hedging is a well-understood topic, we focus our discussion on the use of LBs to hedge longevity risk, and take it as given that the hedger will take account of its effect on its interest-rate risk exposure in designing its interest-rate hedging strategy.

6.2 Elasticities for standard longevity bonds

Figure 4 shows the corresponding elasticity plots for standard LBs. The signs and directions of change of these plots are the same as those of Fig. 3: this makes intuitive sense and suggests that standard LBs have similar qualitative hedging properties as their zero equivalents. However, the actual shapes of the curves are (usually) a little different: in particular, the curves now show some tendency towards flattening out as

³⁰ However, because the sensitivity of the LB depends on the maturity, the amount of the bond required to put on a hedge will vary with the maturity: the longer the maturity, the greater the leverage and the smaller the amount of bond required.

T gets large. The interest-rate elasticities shown in Fig. 4 can be compared with the much greater interest-rate elasticities of comparable conventional annuity bonds shown in Fig. 5.

Insert Figure 4 here

Insert Figure 5 here

6.3 Elasticities for standard floating-coupon principal-at-risk longevity bonds

In contrast, the elasticity curves for the floating-coupon (LIBOR+spread) principal-at-risk LBs shown in Figure 6 is quite different. The q elasticity curve starts at 0, dips and then returns back to 0;³¹ whereas the r elasticity curve also starts at 0, but then rises, peaks and falls back again.³² These securities are long longevity (as well as interest-rate) risk and would therefore be a candidate hedge instrument for a position that is short longevity risk.

Insert Figure 6 here

6.4 Elasticities for standard fixed-coupon principal-at-risk longevity bonds

Finally, the elasticity curves for fixed-coupon principal-at-risk LBs are given in Figure 7. We now find that the switch from floating to fixed coupons makes no difference to the q curve, but makes a big difference to the r curve: this is now negative and falling (instead of positive and peaked). The elasticities have the same signs as their zero and standard equivalents, and therefore have qualitatively (though not quantitatively) similar hedging properties.

³¹ Both the floating- and fixed-coupon principal-at-risk LBs exhibit this same ‘dip’. The explanation for this dip lies in two offsetting effects: as maturity initially rises, any given shock to mortality has a cumulative effect on S_t and the elasticity becomes more pronounced; however, as maturity continues to rise, the elasticity has to fall back again because more and more of the reference population has died off and S_t must eventually approach 0.

³² Again two offsetting effects operate. On the one hand, long maturing bonds have greater interest rate sensitivity than short maturing bonds. On the other hand, the reduced payouts S_t as maturity increases first attenuate and then reverse the first effect.

Insert Figure 7 here

7. Conclusions

Longevity bonds offer a challenging but also very promising new frontier for financial markets. They are challenging because they give rise to issues that are not present with conventional fixed-income securities: their valuation is problematic because of market incompleteness, the construction of longevity bonds is hampered by the absence of ultra-long conventional debt instruments, there can be problems with the survivor indices on which these securities are predicated, and there can be nontrivial problems of contract design (i.e., what particular features are required to ensure that a LB issue is well received in the marketplace?). On the other hand, they also have great promise for institutions wishing to trade exposures to longevity risk. They offer annuity providers and pension funds highly flexible means of managing their exposures to both longevity and interest-rate risk, and we can envisage a great variety of different types of LB, each of these has its own distinct risk management features: its own combination of longevity- and interest-rate risk exposures, its own leverage properties, and so on. At the same time, LBs also have the potential to be very attractive to capital markets institutions in their never-ending search for low-beta investment opportunities – and, indeed, the indications are that these institutions are dying to get involved.

Acknowledgements

The authors thank the UK Economic and Social Research Council for financial support under research grants RES-000-23-1036 and RES-000-27-0014. They also thank participants at Second International Longevity Risk and Capital Markets Solution Conference (Longevity Two) in Chicago on 24 April 2006, especially Paul Hattori, for their helpful comments on an earlier draft. The usual disclaimer applies.

References

- Benjamin, B., and Pollard, J. H. (1993) *The Analysis of Mortality and Other Actuarial Statistics*. 3rd edition. London: Institute of Actuaries.
- Blake, D. (2003) A Reply to ‘Survivor bonds: A comment on Blake and Burrows’. *Journal of Risk and Insurance*, 70, 349-51.

- Blake, D., and Burrows, W. (2001) Survivor bonds: Helping to hedge mortality risk. *Journal of Risk and Insurance*, 68, 339-48.
- Blake, D., Cairns, A.J.G. and Dowd, K. (2006) Living with mortality: Longevity bonds and other mortality-linked securities (with discussion). *British Actuarial Journal*, forthcoming.
- Brown, J. R. and Orszag, P. R. The Political Economy of Government Issued Longevity Bonds. *Journal of Risk and Insurance* 73(4): **This issue**, 2006.
- Cairns, A.J.G. (2004) *Interest-Rate Models: An Introduction*. Princeton: Princeton University Press.
- Cairns, A.J.G., Blake, D., and Dowd, K. (2006a) Pricing death: Frameworks for the valuation and securitisation of mortality risk. *ASTIN Bulletin*, 36, 79-120.
- Cairns, A. J. G., Blake, D. and Dowd, K. (2006b). A Two-Factor Model for Stochastic Mortality with Parameter Uncertainty: Theory and Calibration. *Journal of Risk and Insurance* 73(4): **This issue**, 2006.
- Cowley, A., and Cummins, J.D. (2005) Securitization of life insurance assets and liabilities. *Journal of Risk and Insurance* 72, 193-226.
- Cox, S.H., and Lin, Y. (2004) Natural hedging of life and annuity mortality risks. In *Proceedings of the 14th International AFIR Colloquium*, Boston, 483-507.
- Dahl, M. (2004) Stochastic mortality in life insurance: Market reserves and mortality-linked insurance contracts. *Insurance: Mathematics and Economics*, 35, 113-136.
- Dahl, M., and Møller, T. (2005) Valuation and hedging of life insurance liabilities with systematic mortality risk. In *Proceedings of the 15th International AFIR Colloquium*, Zurich. (Available online at <http://www.afir2005.ch>).
- Dowd, K. (2003) Survivor bonds: A Comment on Blake and Burrows. *Journal of Risk and Insurance*, 70, 339-48.
- Dowd, K., Cairns, A.J.G., and Blake, D. (2006a) Mortality-dependent measures of financial risk. *Insurance: Mathematics and Economics*, 38, 427-440.
- Dowd, K., Blake, D., Cairns, A.J.G., and Dawson, P. E. (2006b) Survivor swaps. *Journal of Risk and Insurance*, 73, 1-17.

Froot, K.A., and Stein, J.C. (1998) Risk management, capital budgeting, and capital structure policy for financial institutions: an integrated approach. *Journal of Financial Economics*, 47, 55-82.

Hardy, M. (2005) A matter of life and death. *Financial Engineering News* No. 46, November-December 2005: 17, 20.

Krutov, A. (2006) Insurance-linked securities: an emerging class of financial instruments. *Financial Engineering News* No. 48, March-April 2006, pp. 7, 16.

Lane, M (2006) Pricing life securitizations and their place in optimal ILS portfolios, presentation at the Second International Longevity Risk and Capital Markets Solutions Conference (Longevity Two), Chicago, Ill, 24 April 2006

Lin, Y., and S. H. Cox (2005) Securitization of mortality risks in life annuities. *Journal of Risk and Insurance* 72: 227-252.

MacMinn, R. and A. Richter (2004). Hedging brevity risk with mortality-based securities, mimeo.

Milevsky, M.A., and Promislow, S.D. (2001) Mortality derivatives and the option to annuitise. *Insurance: Mathematics and Economics*, 29, 299-318.

Miltersen, K.R., and Persson, S.-A. (2005) Is mortality dead? Stochastic forward force of mortality determined by no arbitrage. Working paper, University of Bergen.

Pensions Commission (2005) *A New Pensions Settlement for the Twenty-First Century*, The Stationary Office, Norwich

Perks, W. (1932) On some experiments in the graduation of mortality statistics. *Journal of the Institute of Actuaries* 63: 12-57.

Stallard, E. Demographic Issues in Longevity Risk Analysis. *Journal of Risk and Insurance* 73(4): **This issue**, 2006.

Turner, A. Pensions, Risks and Capital Markets. *Journal of Risk and Insurance* 73(4): **This issue**, 2006.

UK Debt Management Office (2004) *Issuance of Ultra-Long Gilt Instruments: Consultation Document*. London: Debt Management Office, 2 December.

Wang, S.S. (1996) Premium calculation by transforming the layer premium density. *ASTIN Bulletin*, 26, 71-92.

Wang, S.S. (2000) A class of distortion operations for pricing financial and insurance risks. *Journal of Risk and Insurance*, 67, 15-36.

Wang, S.S. (2002) A universal framework for pricing financial and insurance risks. *ASTIN Bulletin*, 32, 213-234.

Wang, S.S. (2003) Equilibrium pricing transforms: new results using Buhlmann's 1980 economic model. *ASTIN Bulletin*, 33, 57-73.

Appendix A: The Mortality Model Used in Section 6

The mortality forecasting model used in section 6 is that of Cairns *et alia*, 2006b, and works as follows. Let $S_{t,x}$ be the survivor rate at time t of a cohort aged x in year 0. We know that if $q_{t,x}$ is the realized mortality rate in year $t+1$ (that is, from time t to time $t+1$) of our cohort, then

$$S_{t+1,x} = (1 - q_{t,x})S_{t,x} \quad (\text{A1})$$

We now assume that $q_{t,x}$ is governed by the following two-factor Perks stochastic process:

$$q_{t,x} = e^{A_1(t+1) + A_2(t+1, (t+x))} / (1 + e^{A_1(t+1) + A_2(t+1, (t+x))}) \quad (\text{A2})$$

where $A_1(t+1)$ and $A_2(t+1)$ are stochastic processes measurable at time $t+1$ (see Perks, 1932, Benjamin and Pollard, 1993). Cairns *et alia*, 2006b generate empirical results showing that this mortality model provides a good fit to realized male mortality data in England and Wales. Their results also indicate that a two-factor model of UK mortality fits the data better than a one-factor one.

Now let $A(t) = (A_1(t), A_2(t))'$ and assume that $A(t)$ is a random walk with drift:

$$A(t+1) = A(t) + b + CZ(t+1) \quad (\text{A3})$$

where b is a constant 2×1 vector of drift parameters, C is a constant 2×2 lower triangular matrix reflecting volatilities and correlations, and $Z(t+1)$ is a 2×1 vector of independent standard normal variables.³³ Cairns *et alia*, 2006b also show

³³ This model also involves an additional assumption that interest rates evolve independently of mortality over time. There is anecdotal evidence (see, for example, Miltersen and Persson, 2005) that

that if we use the UK Government Actuary's Department (GAD) data for English and Welsh males over 1961-2002, then the least squares estimates of our parameters are:

$$\hat{b} = \begin{bmatrix} -0.04340 \\ 0.000367 \end{bmatrix} \quad (\text{A4a})$$

$$\hat{\Sigma} = \hat{C}\hat{C}' = \begin{bmatrix} 0.01067000, & -0.00016170 \\ -0.00016170, & 0.00000259 \end{bmatrix} \quad (\text{A4b})$$

We can recover \hat{C} from $\hat{\Sigma}$ using a Choleski decomposition, and all that remains is to specify a suitable starting value $A(0)$. The results of Cairns *et alia*, 2006b, Fig. 2, suggest that we might take $A(0) \approx (-11.0, 0.107)'$ if we take 2003 as our starting point (i.e., if we set $t=0$ for the end of 2003).

Having specified the model, we simulate paths for $A(t)$ over each of $t=1,2,\dots,50$, using our assumed values of $A(0)$. Each path of $A(t)$ values gives us a path of realized mortality rates $q_{t,x}$, and each such path gives us a path for the survivor rates $S_{t,x}$.

this assumption might not be valid over long periods of time, in which case the pricing formula above needs to be adjusted to reflect the linkage between the interest-rates and mortality. This can be done by using, for example, forward-pricing measures rather than the risk-neutral measure (see, for example, Cairns, 2004).

Appendix B: Sensitivities and Hedging Uses of the Inverses of the Longevity Bonds Considered in Section 6

Plots of the elasticities for inverse LZs are given in Figure B.1.³⁴ The q elasticity starts at 1 and then declines smoothly towards 0 as T gets large. This implies that an inverse LZ can only be a useful longevity hedge if it has a short to medium maturity, since at long maturities, $1 - S_T$ approaches unity. For its part, the r elasticity of the inverse LZ is the same as that of a LZ of the same maturity. The curves in Fig. B.1 therefore indicate the holder of an inverse LZ is short both longevity and interest rate risk. This suggests that a long position in an inverse LZ is a hedge for a position that is long longevity risk (e.g., as might be the case for a provider of life products), and a corresponding short position might be a suitable hedging instrument for a position that is short longevity risk (e.g., as might be the case for an annuity provider). The elasticities for inverse LBs are given in Figure B.2. They are similar to those of inverse LZs, but are more attenuated as their maturity increases.

Insert Figure B.1 here

Insert Figure B.2 here

In contrast, the elasticity curves for inverse floating-coupon principal-at-risk LBs shown in Figure B.3 are quite different. The q elasticity is much lower at all maturities³⁵, while the r elasticity is non-negative at all maturities. This means that the inverse security is short longevity risk but long interest rate risk, suggesting that it might hedge a position that was long longevity risk. The elasticities of inverse fixed-coupon principal-at-risk LBs are shown in Figure B.4. The q elasticity is very

³⁴ The explanations for the shapes of these (and the later) curves are intuitively fairly obvious and do not need spelling out; instead, we prefer to focus on their hedging implications.

³⁵ There is no necessary ‘dip’ with inverse principal-at-risk LBs because although these bonds have elasticities that also tend to zero as maturity gets large, their elasticities (unlike those of standard principal-at-risk bonds) start at significantly positive values and then decline, whereas standard principal-at-risk bonds ‘dip’ because their elasticities start at 0 as well as end at 0.

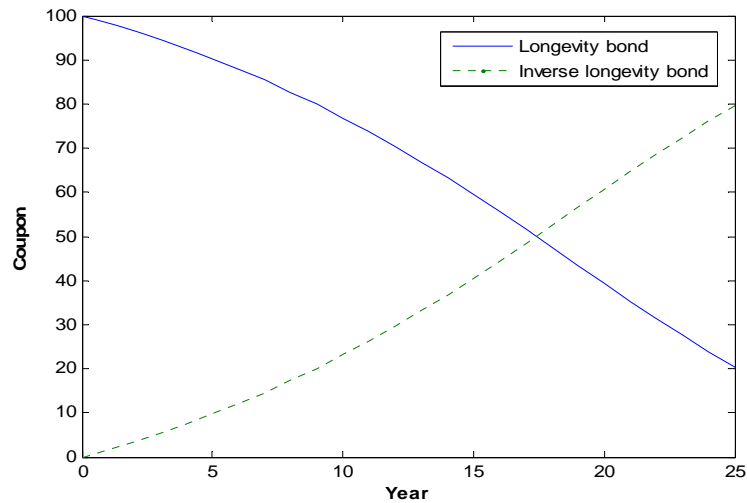
similar to the floating-coupon case, but the r elasticity is now negative at all maturities. The longevity risk hedging properties will therefore be very similar to those of the floating coupon bond, but there will be an inverse response to interest rate risk.

Insert Figure B.3 here

Insert Figure B.4 here

FIGURES

Figure 1. Coupon Payments on a 25-year Longevity Bond and Inverse Longevity Bond Constructed from Decomposing an Annuity Bond



Note: Estimating using 5000 simulation trials of the Cairns *et alia*, 2006b, mortality forecasting model outlined in the Appendix. The model is calibrating for 65-year old males using data for English and Welsh males over the period 1961-2002 provided by the UK Government Actuary's Department. The initial parameter values $A_1(0)$ and $A_2(0)$ were taken as -11 and 0.107 as in Cairns *et alia*, 2006b, Figure 2. The LB pays S_t and the inverse LB pays $1 - S_t$ in period t .

Figure 2. A Longevity Bond Constructed from Longevity Zeros and a Longevity Swap

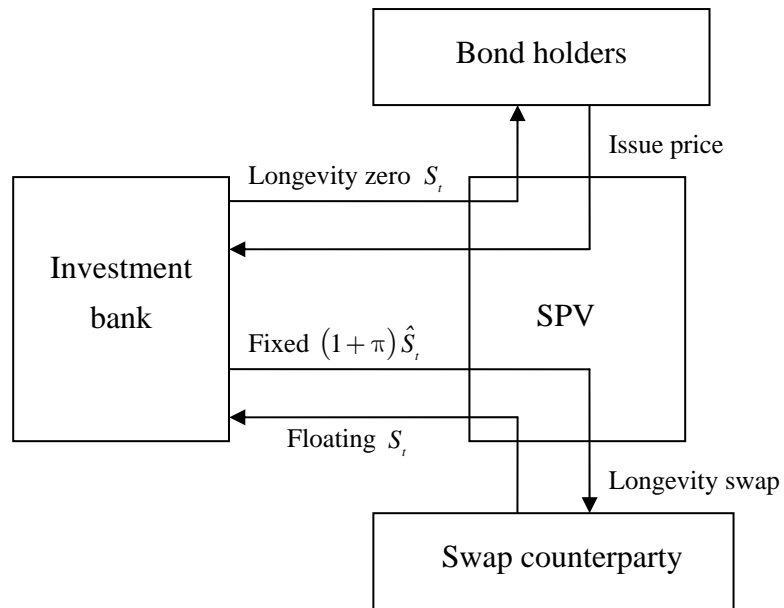




Figure 3. Elasticities for Longevity Zeros




Note: The longevity zero pays S_T in period T . The parameters q and r are the mortality rate and the interest rate, respectively. Estimated using 5000 simulation trials of the Cairns *et alia*, 2006b, mortality forecasting model outlined in Appendix A. The model is calibrated for 65-year old males using data for English and Welsh males over the period 1961-2002 provided by the UK Government Actuary's Department. The initial parameter values $A_1(0)$ and $A_2(0)$ were taken as -11 and 0.107 as in Cairns *et alia*, 2006b, Figure 2, the discount rate was taken as 0.05, and the values of the market prices of risk λ_1 and λ_2 were taken as 0.175.

Figure 4. Elasticities for Standard Longevity Bonds




Note: As per the note to Figure 3, except that the standard LB pays S_t for each period t between 1 and T .

Figure 5. Interest-rate Elasticity for Conventional Annuity Bonds




Note: The annuity bond pays a coupon of one unit for each period t from 1 to $T-1$.

Figure 6. Elasticities for Floating-Coupon Principal-at-Risk Longevity Bonds




Note: As per the note to Figure 3, except that the LB pays a principal of S_T in period T and coupon payments are floating and equal to the market interest rate.

Figure 7. Elasticities for Fixed-Coupon Principal-at-Risk Longevity Bonds




Notes: As per the notes to Figure 6, except that coupon payments are set at the initial discount rate.

Figure B.1. Elasticities for Inverse Longevity Zeros




Note: The inverse longevity zero pays $1 - S_T$ in period T . The parameters q and r are the mortality rate and the interest rate, respectively. Estimated using 5000 simulation trials of the Cairns *et alia*, 2006b, mortality forecasting model outlined in Appendix A. The model is calibrated for 65-year old males using data for English and Welsh males over the period 1961-2002 provided by the UK Government Actuary's Department. The initial parameter values $A_1(0)$ and $A_2(0)$ were taken as -11 and 0.107 as in Cairns *et alia*, 2006b, Figure 2, the discount rate was taken as 0.05, and the values of the market prices of risk λ_1 and λ_2 were taken as 0.175.

Figure B.2. Elasticities for Inverse Longevity Bonds




Note: As per the note to Figure B.1, except that the inverse LB pays $1 - S_t$ for each period t between 1 and T .

Figure B.3. Elasticities for Floating-Coupon Principal-at-Risk Inverse Longevity Bonds



Note: As per the note to Figure B.1 except that the bond pays a principal of $1 - S_T$ in period T . Coupon payments are floating and equal to the market interest rate.

Figure B.4. Elasticities for Fixed-Coupon Principal-at-Risk Inverse Longevity Bonds



Note: As per the note to Figure B.3, except that coupon payments are set at the initial discount rate.