

LIVING WITH MORTALITY: LONGEVITY BONDS AND OTHER MORTALITY-LINKED SECURITIES

BY D. BLAKE, A. J. G. CAIRNS AND K. DOWD

[Presented to the Faculty of Actuaries, 16 January 2006]

ABSTRACT

This paper addresses the problem of longevity risk — the risk of uncertain aggregate mortality — and discusses the ways in which life assurers, annuity providers and pension plans can manage their exposure to this risk. In particular, it focuses on how they can use mortality-linked securities and over-the-counter contracts — some existing and others still hypothetical — to manage their longevity risk exposures. It provides a detailed analysis of two such securities — the Swiss Re mortality bond issued in December 2003 and the EIB/BNP longevity bond announced in November 2004. It then looks at the universe of hypothetical mortality-linked securities — other forms of longevity bonds, swaps, futures and options — and investigates their potential uses. It also addresses implementation issues, and draws lessons from the experiences of other derivative contracts. Particular attention is paid to the issues involved with the construction and use of mortality indices, the management of the associated credit risks, and possible barriers to the development of markets for these securities. It suggests that these implementation difficulties are essentially teething problems that will be resolved over time, and so leave the way open to the development of flourishing markets in a brand new class of securities.

KEYWORDS

Longevity Risk; Longevity Bonds; Mortality Swaps; Annuity Futures; Mortality Options; Mortality Swaptions; Survivor Caps; Mortality Indices; Liquidity; Basis Risk; Credit Risk

CONTACT ADDRESSES

David Blake, Pensions Institute, Cass Business School, 106 Bunhill Row, London EC1Y 8TZ, U.K. Tel:+44 (0) 20 7040 5143; e-mail: D.Blake@city.ac.uk

Andrew Cairns, Actuarial Mathematics and Statistics, School of Mathematical and Computer Sciences, Heriot-Watt University, Edinburgh EH14 4AS, U.K. Tel:+44 (0) 131 451 3245; e-mail: A.Cairns@ma.hw.ac.uk

Kevin Dowd, Centre for Risk and Insurance Studies, Nottingham University Business School, Jubilee Campus, Nottingham NG8 1BB, U.K. Tel:+44 (0) 115 846 6682; e-mail: Kevin.Dowd@nottingham.ac.uk

1. INTRODUCTION

1.1 *Background*

1.1.1 Benjamin Franklin once famously observed that nothing is certain

in life except death and taxes. Death presumably remains as inevitable as it always was, but over the last few decades it has also become clear that the timing of death is getting later on average. When the British welfare state began in 1948, men had a life expectancy of 67 years. By the beginning of the 21st century, many British men now live into their early 80s, and women can expect to live into their late 80s. It has also been evident for many years that mortality rates have been evolving in an apparently stochastic fashion. This phenomenon — the uncertainty attached to aggregate mortality rates — has been given the name longevity risk.

1.1.2 Longevity risk is a key risk factor in many life assurance and pensions products. For example, annuity providers are exposed to the risk that the mortality rates of pensioners will fall at a faster rate than anticipated in their pricing and reserving calculations. Yet annuities are commoditised products. They sell mainly on the basis of price (although factors such as service and credit rating are also considerations), and profit margins have to be kept low to gain and then protect market share. If the mortality assumptions built into the prices of annuities turn out to be overestimates, this cuts straight into the profit margin of annuity providers. Indeed, some life companies claim to be losing money on their annuity business because annuitants already live too long; as a result some companies now cover themselves against the risk of further longevity improvements by only quoting on uncompetitive terms (see Section 8.4, Figure 5).

1.1.3 Annuities in their various forms (see, for example, Wadsworth, Findlater & Boardman, 2001) are the mainstay of pension plans throughout the world, and this is especially the case in the United Kingdom. They are the only retail instrument ever devised capable of hedging longevity risk. Without them, pension plans would have great difficulty performing their fundamental task of protecting retirees from outliving their resources for however long they live.

1.1.4 The possible consequences of longevity risk came to public prominence in December 2000 when the world's oldest life office, the Equitable Life Assurance Society (ELAS), was forced to close to new business. Between 1957 and 1988, ELAS had sold with-profit pension annuities with 'guaranteed annuity rates' (GARs) fixed by reference to specific assumptions regarding interest rates and life expectancy. These embedded options became very valuable in the 1990s due to a combination of falling interest rates and improvements in mortality, and it was the rise in the values of these options that led to ELAS's financial difficulties. These could have been avoided if ELAS had hedged its exposure to both interest-rate risk and longevity risk, but for years ELAS failed to appreciate the extent of its potential exposure. The failure of ELAS to do so bespeaks of the poor state of interest-rate and longevity risk management in the Society. However, even if it had anticipated the problem, it still

lacked good instruments to hedge its exposure to both risks, particularly longevity.

1.2 *Focus of this Paper*

1.2.1 This paper addresses the issue of how to manage (aggregate or systematic) longevity risk. We do not consider non-systematic longevity risk (or the risk associated with the status of individual lives).

1.2.2 Nor do we consider detailed approaches to the modelling of longevity risk in this paper. This topic is treated in detail elsewhere. For example, Cairns, Blake & Dowd (2005b) describe a range of frameworks that can be employed in stochastic modelling (an earlier version of this appears in the proceedings of the 2004 AFIR Colloquium — see Cairns, Blake & Dowd, 2004).

1.2.3 Instead, we focus here on the ways in which life insurers and pension plans can manage their exposure to longevity risk and, most especially, on the ways in which they can use mortality-linked securities — some already existing and others still hypothetical — to manage their longevity risk exposures.

1.3 *Layout of this Paper*

The paper is organised as follows. Section 2 discusses the range of possible ways in which life companies and pension plans might manage longevity risk, and focuses on how they might manage this risk using mortality-linked securities. Section 3 then discusses the different stakeholders in the markets for mortality-linked securities. Section 4 examines the limited range of such securities currently available. Section 5 discusses the much broader range of hypothetical mortality-linked securities — bonds, swaps, futures and options — and Sections 6 to 9 examine each of these in turn and consider how they might be used. Section 10 then addresses the important question of the construction and use of the underlying mortality index that would determine the payments to be made. Section 11 examines the credit risk issues that can arise with mortality-linked securities. Section 12 discusses possible barriers to the development of healthy markets in these securities, and Section 13 concludes.

2. A RANGE OF POSSIBLE RESPONSES

2.1 Companies affected by longevity risk can respond to it in a number of ways:

- (i) They can simply accept the risk as a legitimate business risk that they understand well and are prepared to assume.
- (ii) They can diversify their longevity risks (e.g. across different products, countries or socio-economic groups). Similarly, they can balance their

portfolio by seeking to exploit possible natural hedges involved running a mixed business of term assurance and annuity business (see, for example, Cox & Lin, 2004).

- (iii) They can enter into a variety of forms of reinsurance with a reinsurance company. Such contracts might involve sharing some or all of the downside of longevity risk with the reinsurer.
- (iv) Pension plans can arrange for a bulk buy out of their pensions in payment, passing on responsibility for payment to an insurance company.
- (v) Smaller pension plans are accustomed to purchasing annuities at the time of retirement for each plan member. This hedges the total risk (both systematic and non-systematic) in their pool of pensioners. Unless the plan purchases deferred pensions on a regular basis, it still bears the longevity risk for current active members and deferred pensioners between now and their retirement dates.
- (vi) Annuity providers might choose to replace traditional non-participating annuities with participating contracts that pass part of the exposure to longevity risk on to the surviving participating policyholders. For example, Wadsworth, Findlater & Boardman (2001) and Blake, Cairns & Dowd (2003) both describe mechanisms that pay bonuses or survivor credits to annuitants that can take account of experienced mortality rates within the pool of annuitants. This is in contrast to the traditional annuity under which survivor credits are based on mortality rates predicted at the time of purchase.
- (vii) They might securitise a line of business with a high level of longevity risk. (For further discussion of the issues involved with securitisation, see Cowley & Cummins, 2005.)
- (viii) They can manage the risk using mortality-linked securities. These securities might be traded contracts (such as longevity bonds, annuity futures, options, etc.) or over-the-counter (OTC) contracts (such as mortality swaps or forwards). (An OTC contract is not strictly a 'security'. However, for the sake of brevity we will assume that the expression mortality-linked securities includes OTC contracts.)

2.2 Our focus of interest in this paper is with (viii), the use of mortality-linked securities to manage longevity risk, based on the underlying assumption — which we believe will often be a reasonable one — that the main parties concerned wish to hedge this risk.

3. STAKEHOLDERS IN MARKETS FOR MORTALITY-LINKED SECURITIES

3.1 *Classes of Stakeholders*

Before examining these securities, it is helpful to discuss who might be

interested in the markets for mortality-linked securities. These markets have a number of stakeholders.

3.2 *Hedgers*

One natural class of stakeholders are hedgers, those who have a particular exposure to longevity risk and wish to lay off that risk. For example, annuity providers stand to lose if mortality improves by more than anticipated, whilst life assurers stand to gain, and vice versa. These offsetting exposures imply that annuity providers and life assurers can hedge each other's longevity risks. (In many cases, annuity providers and life assurers are part of the same life office, in which case the annuity and life books provide at least a partial natural hedge.) Alternatively, parties with unwanted exposure to longevity risk might pay other parties to lay off some of their risk. For instance, a life office might hedge its longevity risk by reinsuring it, or by transferring it to the capital markets.

3.3 *General Investors*

Provided expected returns are reasonable, capital market institutions such as investment banks or hedge funds might be interested in acquiring an exposure to longevity risk, since it has a low correlation with standard financial market risk factors. The combination of a low beta and a potentially positive alpha should therefore make mortality-linked securities attractive investments in diversified portfolios.

3.4 *Speculators and Arbitrageurs*

A market in mortality-linked securities might attract speculators: short-term investors who trade their views on the direction of individual security price movements. The active involvement of speculators is very helpful for market liquidity, and is in fact essential to the success of traded futures and options markets (see Sections 8.2 and 8.3). Arbitrageurs seek to profit from any pricing anomalies in related securities. For arbitrage to be a successful activity, it is essential that there are well-established pricing relationships between the related securities.

3.5 *Government*

3.5.1 The Government has many potential reasons to be interested in markets for mortality-linked securities. It might wish to promote such markets and assist financial institutions that are exposed to longevity risk (e.g. it might issue longevity bonds that can be used as instruments to hedge longevity risk — see Section 6). Actions of this type potentially reduce the probability that large companies are bankrupted by their pension funds, with the result that society as a whole benefits from the greater stability of the economy. As 'insurer of last resort', the Government is also potentially the residual holder of this risk in the event of default by private sector pension

funds and insurance companies. In the U.K., for example, the Government has a strong incentive to help companies hedge their exposure to longevity risk, which would reduce the likelihood of claims on the new Pension Protection Fund.

3.5.2 The Government might also be interested in managing its own exposure to longevity risk. The Government is a significant holder of this risk in its own right, via the pay-as-you-go state pension system, via its obligations to provide health care for the elderly, and for many other similar reasons. At a higher level, the Government is affected by numerous other economic factors, some of which partially offset the Government's own exposure to longevity risk (for example, income tax payable on private pensions in payment).

3.6 *Regulators*

Financial regulators have two main stated aims: (i) the enhancement of financial stability through the promotion of efficient, orderly and fair markets, and (ii) ensuring that retail customers get a fair deal. (These are the stated aims of the U.K. Financial Services Authority as set out in the Financial Services and Markets Act 2000.)

3.7 *Other Stakeholders*

Other stakeholders might include security managers and organised exchanges, both of which would benefit from a new source of fee income.

4. EXISTING MORTALITY-LINKED SECURITIES

4.1 *Existing Securities*

We will now describe the main mortality-linked securities that currently exist and/or have been announced. These are the Swiss Re mortality bond and the EIB/BNP longevity bond.

4.2 *The Swiss Re Mortality Bond*

4.2.1 In December 2003, Swiss Re issued a three-year life catastrophe bond, maturing on 1 January 2007, which helps to reduce Swiss Re's exposure to a catastrophic mortality deterioration (e.g. such as that associated with a repeat of the 1918 Spanish Flu pandemic). The issue size was \$400m. Investors receive quarterly coupons set at three-month U.S. dollar LIBOR + 135 basis points.

4.2.2 However, the principal is unprotected and depends on what happens to a specifically constructed index of mortality rates across five countries: the United States of America, U.K., France, Italy and Switzerland. The principal is repayable in full if the mortality index does not exceed 1.3 times the 2002 base level during any of the three years of the

Table 1. Swiss Re mortality bond payoff schedule

Payment at maturity (T)	$100\% - \sum_t loss_t$	if $\sum_t loss_t < 100\%$
	0%	if $\sum_t loss_t \geq 100\%$
Loss percentage in year $t = loss_t$	0%	if $q_t < 1.3q_0$
	$[(q_t - 1.3q_0)/(0.2q_0)] \times 100\%$	if $1.3q_0 \leq q_t \leq 1.5q_0$
	100%	if $1.5q_0 < q_t$

where:

$q_0 =$ base index
 $q_t = \sum_j C_j \sum_i (G^m A_i q_{i,j,t}^m + G^f A_i q_{i,j,t}^f)$

Key :

$q_{i,j,t}^m =$ mortality rate (deaths per 100,000) for males in the age group i for country j
 $q_{i,j,t}^f =$ mortality rate (deaths per 100,000) for females in the age group i for country j
 $C_j =$ weight attached to country j
 $A_i =$ weight attributed to age group i (same for males and females)
 G^m and $G^f =$ gender weights applied to males and females respectively

The following country weights apply:

U.S.A. 70%, U.K. 15%, France 7.5%, Italy 5%, Switzerland 2.5%, male 65%, female 35%.

bond’s life. The principal is reduced by 5% for every 0.01 increase in the mortality index above this threshold and is completely exhausted if the index exceeds 1.5 times the base level. The payoff schedule of the bond is shown in Table 1 and Figure 1.

4.2.3 The bond was issued via a special purpose vehicle (SPV) called Vita Capital (VC). VC invests the \$400m principal in high-quality bonds and

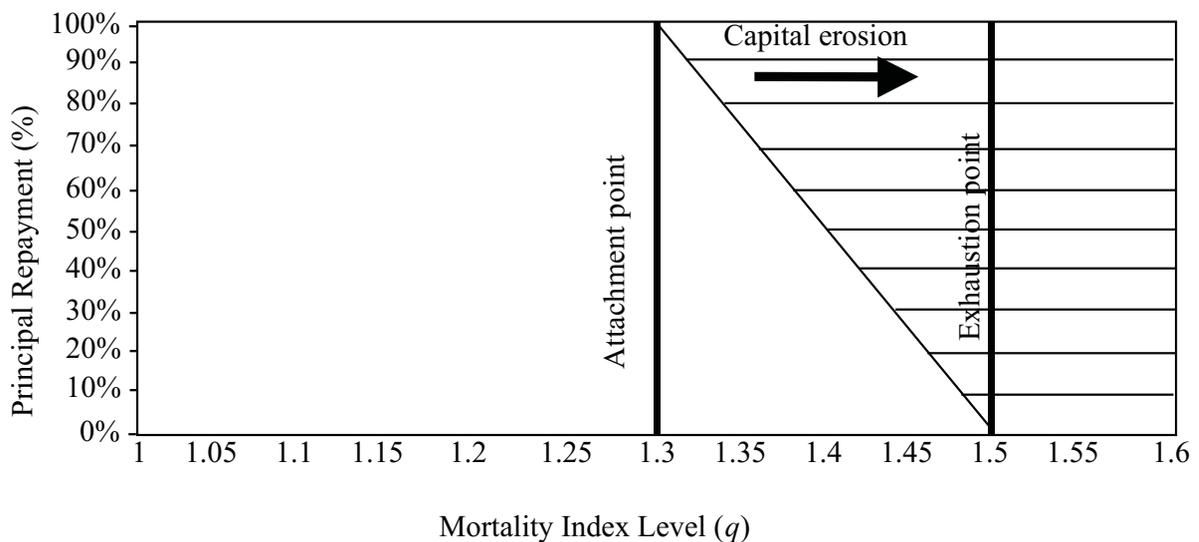


Figure 1. Terminal payoff of Swiss Re mortality bond to investors

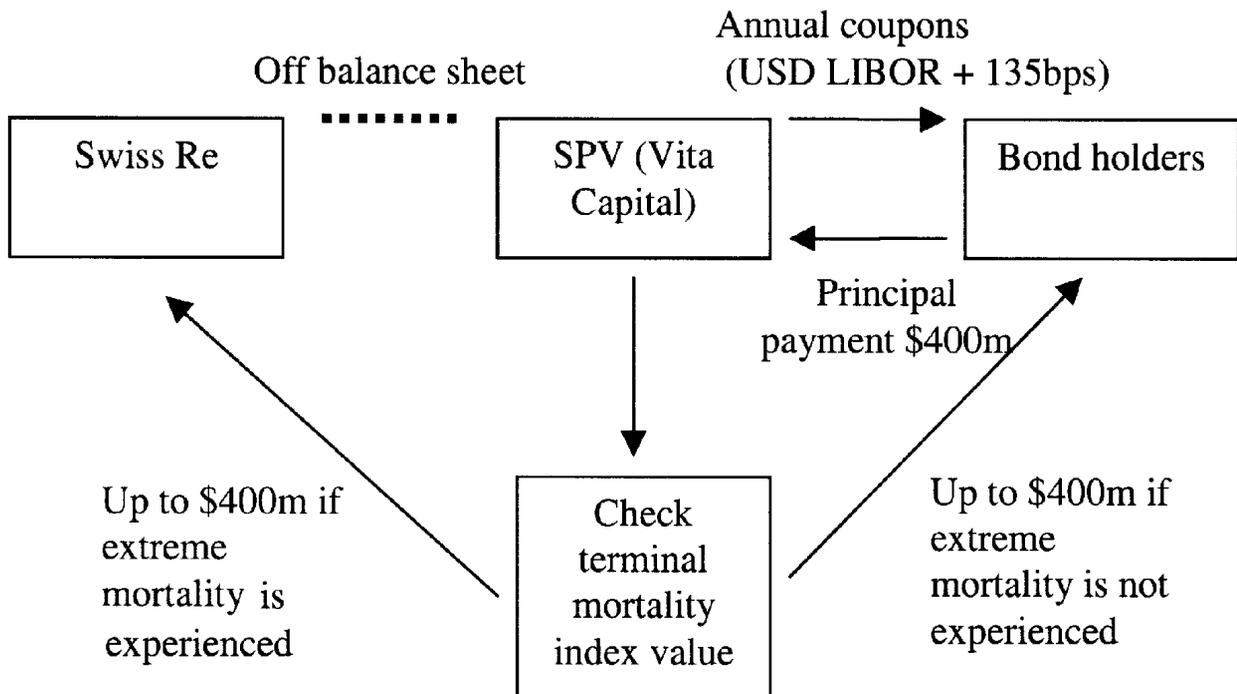


Figure 2. The structure of Swiss Re mortality bond

swaps the income stream on these for a LIBOR-linked cash flow. VC distributes the quarterly income to investors and any principal repayment at maturity. This structure is shown in Figure 2. The benefits of using an SPV in this context are that the cash flows are kept off balance sheet (which is good from Swiss Re's point of view) and that credit risk is reduced (which is good from the investors' point of view).

4.2.4 According to its 2004 annual report, life reinsurance is Swiss Re's primary source of business revenue, accounting for 30% of revenues, implying that profitability is negatively correlated with mortality rates. However, as the world's largest provider of life and health reinsurance, Swiss Re faces the potential difficulty of finding a sufficient number of counterparties on whom it can offload this risk, and this has implications for its regulatory capital requirements. The bond therefore helps Swiss Re to unload some of the extreme mortality risk that it faces. It is also likely that Swiss Re was mindful of its credit rating and wanted to reassure rating agencies about its mortality risk management. Further, by issuing the bond themselves, Swiss Re are not dependent on the creditworthiness of other counterparties should an extreme mortality event occur. Thus, the bond gives Swiss Re some protection against extreme mortality risk without requiring that the company acquire any credit risk exposure in the process.

4.2.5 Investors in the bond take the opposite position and receive an enhanced return if an extreme mortality event does not occur. Some

indication of how well compensated they were for taking on this extreme mortality risk arises from the work of Beelders & Colarossi (2004). They valued the bond using extreme value theory, assuming a generalised Pareto distribution for mortality. Recognising that the terms of the bond are equivalent to a call option spread on the mortality index, with a lower strike price of $1.3q_0$ and an upper strike price of $1.5q_0$, Beelders and Colarossi estimated the value of the probability of attachment ($\text{prob}[q_t > 1.3q_0]$) at 33 basis points and the value of the probability of exhaustion ($\text{prob}[q_t > 1.5q_0]$) at 15 basis points. The expected loss on the bond was estimated to be 22 basis points, less than the 135 basis points of compensation on offer initially to investors. Beelders and Colarossi concluded that the bond appeared to be a good deal for investors and in June 2004 the bond was trading at LIBOR + 100 basis points. However, we should keep in mind that their figures are only estimates based on a model that ignores parameter uncertainty: plausible alternative parameter estimates can produce much higher values for the basis point compensation received by investors. Thus, we cannot be sure how good a deal the investors actually got. By November 2005 the mid-market price of the bond was equivalent to LIBOR + 123 basis points. It is plausible (although we have no evidence for this) that this increase reflected the increased probability of a bird-flu pandemic in 2006.

4.2.6 The Swiss Re bond issue was fully subscribed and press reports suggest that investors were happy with it (e.g. *Euroweek*, 19 December 2003). These investors included a number of pension funds. These would have been attracted, in part, by the higher coupons being offered. They would also have been attracted by the hedging opportunities offered by the fact that the mortality risk associated with the bond is correlated with the mortality risk associated with active members of a pension plan. Specifically, consider an event that would trigger a reduction in the repayment of the Swiss Re bond. The large number of extra deaths would presumably extend to active members of the pension plan. Since death benefits are typically less than the pension liability for an individual member, the reduction in the value of the pension plan's Swiss Re bond investment would be matched by a reduction in the value of their plan liabilities. In the meantime, the bond offers a considerably higher return than similarly rated floating-rate securities. The bond's reception in the marketplace also suggests that investors believed the 135 basis points to represent a good deal.

4.2.7 In April 2005, Swiss Re announced that it had issued a second life catastrophe bond with a principal of \$362m, using a new SPV called Vita Capital II. The maturity date is 2010 and the bond was issued in three tranches: Class B (\$62m), Class C (\$200m) and Class D (\$100m). The principal is at risk if, for any two consecutive years before maturity, the combined mortality index exceeds specified percentages of the expected mortality level (120% for Class B, 115% for Class C, and 110% for Class D). The bond was fully subscribed.

4.3 The EIB/BNP Longevity Bond

4.3.1 In November 2004, BNP Paribas announced a further innovation, a long-term longevity bond targeted at pension plans and other annuity providers. This particular security was not as well received by investors, did not generate sufficient demand to be launched, and was later (late 2005) withdrawn for redesign. However, it received a great deal of public attention and we examine it here in some detail because it is a very interesting security and provides an instructive case study.

4.3.2 This security was to be issued by the European Investment Bank (EIB), with BNP Paribas as the designer and originator and Partner Re as the longevity risk reinsurer. The face value of the issue was £540 million and the bond had a 25-year maturity. The bond was an annuity (or amortising) bond with floating coupon payments, and its innovative feature was to link the coupon payments to a cohort survivor index based on the realised mortality rates of English and Welsh males aged 65 in 2002. The initial coupon was set at £50 million.

4.3.3 In the absence of credit risk, the contract cash flows are simple to specify. For simplicity we will refer to 31 December 2004 as time $t = 0$, with $t = 1$ representing 31 December 2005 etc. Now let $m(y, x)$ represent the crude central death rate for age x published by the Office for National Statistics in the year y . We then construct a survivor index $S(t)$ as follows:

$$S(0) = 1$$

$$S(1) = S(0) \times (1 - m(2003, 65))$$

$$S(t) = S(0) \times (1 - m(2003, 65)) \times (1 - m(2004, 66)) \times \dots \times (1 - m(2002 + t, 64 + t)).$$

At each time $t = 1, 2, \dots, 25$, the bond pays a coupon of £50 million $\times S(t)$.

4.3.4 These cash flows are illustrated in Figure 3. As far as investors are concerned, they make an initial payment of around £540 million (i.e. the

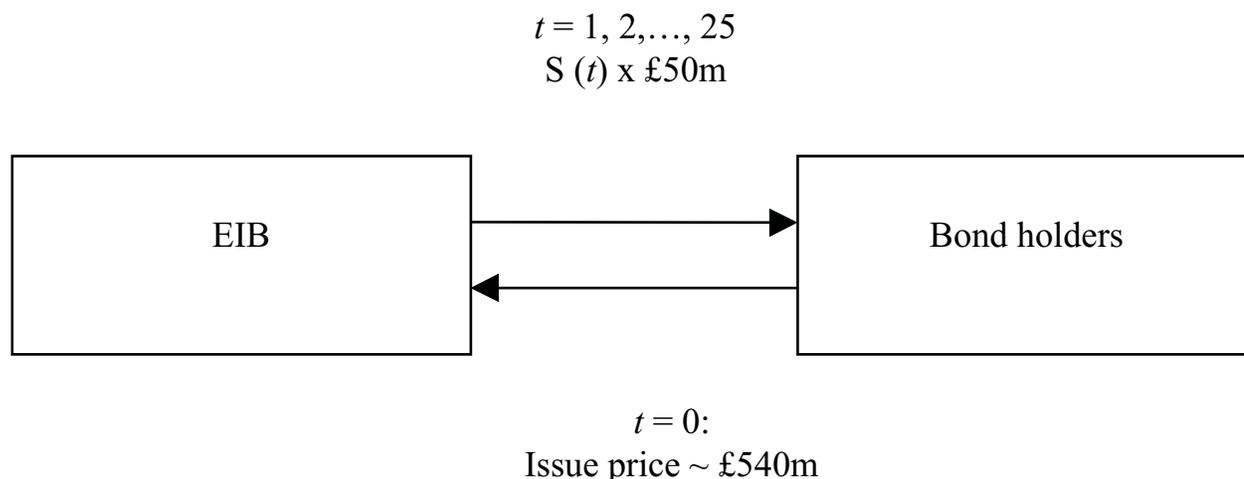


Figure 3. Cash flows from the EIB/BNP Bond, as viewed by investors

issue price) and receive in return an annual mortality-dependent payment of £50 million $\times S(t)$ in each year t for 25 years.

4.3.5 Although the bond was never launched, the issue price was determined by BNP Paribas as follows:

- (i) Ignoring for the moment the £50 million multiplier, the contract specifies a set of anticipated cashflows $\bar{S}(t)$ based on the Government Actuary's Department's 2002-based projections of mortality.
- (ii) Each projected cashflow is priced by discounting at LIBOR minus 35 basis points. The EIB curve typically stands about 15 basis points below the LIBOR curve, so that investors in the longevity bond are being asked to pay 20 basis points to hedge their longevity risk. For further discussion of this risk premium, the reader is referred to Cairns, Blake, Dawson & Dowd (2005) and Cairns, Blake & Dowd (2005a).

4.3.6 The details given above describe the cash flows from the point of view of the investors. However, there are also issues of credit risk to consider, and these lead to some complex background details. These details and the involvement of BNP Paribas and Partner Re are represented in Figure 4. The

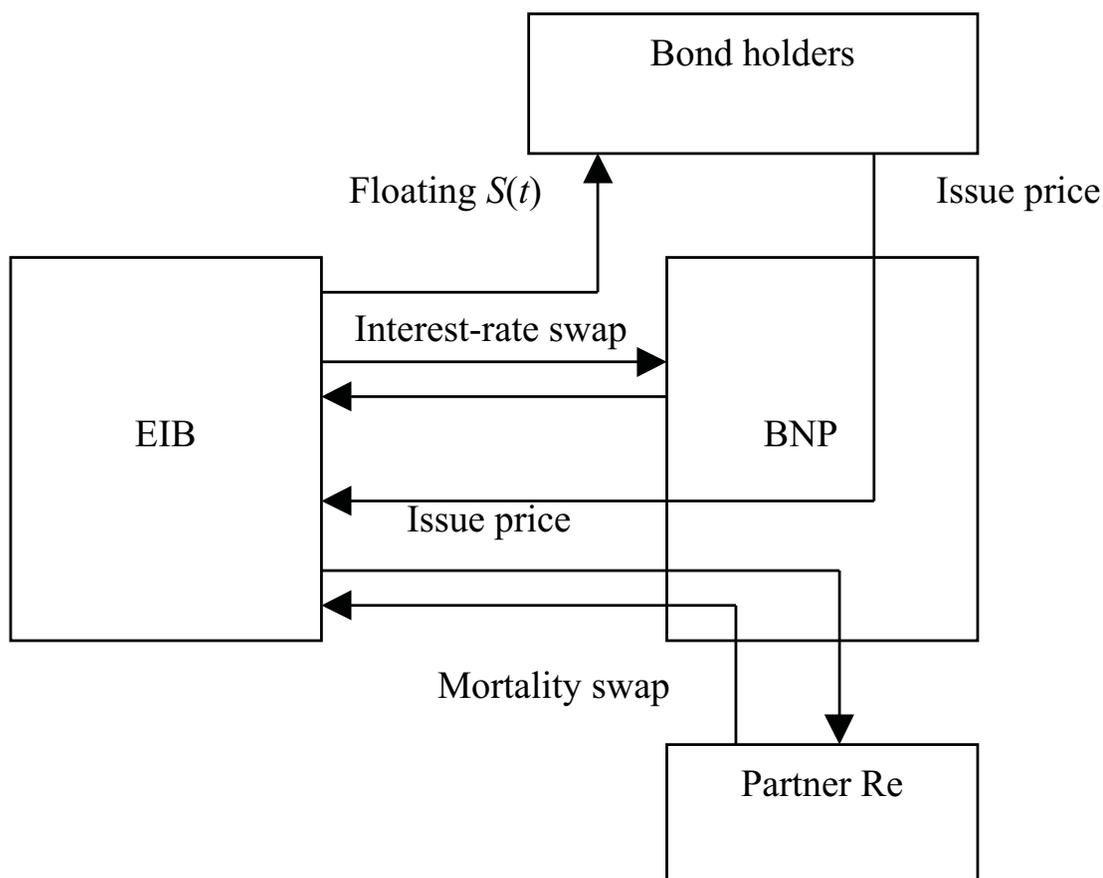


Figure 4. Cash flows from the EIB/BNP bond

longevity bond is actually made up of 3 components. The first is a floating rate annuity bond issued by the EIB with a commitment to pay in euros (€). The second is a cross-currency interest-rate swap between the EIB and BNP Paribas, in which the EIB pays floating euros and receives fixed sterling. These fixed payments, $\hat{S}(t)$, might be, but do not have to be, equal to the $\bar{S}(t)$. (The fixed rate, $\hat{S}(t)$, has to be set to ensure that the swap has zero value at initiation. Typically, this would require the fixed rate to be close but not equal to $\bar{S}(t)$.) From the EIB's perspective, this converts the first element, the floating-rate € bond, into a fixed-rate £ bond. The third and most distinctive component is a mortality swap between the EIB and Partner Re, in which the EIB exchanges the fixed sterling $\hat{S}(t)$ for the floating sterling $S(t)$ at each of the payment dates, $t = 1, \dots, 25$. Strictly speaking, the third component is an OTC deal between BNP and Partner Re. The second component then becomes a commitment from BNP to pay $\text{£}S(t)$ to the EIB, rather than $\text{£}\bar{S}(t)$, in return for floating €. For this reason, we see in Figure 4 that the mortality-swap cash flows are directed through BNP. Ignoring credit risk, the result of the two swaps from the perspective of the EIB is to convert floating € into $\text{£}S(t)$. The intermediate swap of floating € for floating $\text{£}\hat{S}(t)$ does not (as noted above) require that $\hat{S}(t) = \bar{S}(t)$: the price agreed for this swap will, however, depend on what level the $\hat{S}(t)$ are set at. Similarly the price for the mortality swap will depend on the $\hat{S}(t)$.

4.3.7 Note that the second component implies that EIB and BNP have potential credit exposures to each other, and such exposures would become manifest if underlying random factors change and the swap value moves away from 0 (in which case the swap would become an asset to one party and a liability to the other). The third component implies that BNP has a credit exposure to Partner Re. The parties concerned might (or might not) wish to take out some form of insurance on these various credit exposures. The issue of credit risk is discussed further in Section 11.

4.3.8 It is important to appreciate what is going on here in plain language. In a nutshell, the bond is issued by the EIB, and investors only face a credit exposure to the EIB. The EIB has a commitment to make mortality-linked payments in sterling, and then engages in a swap with BNP to exchange this commitment for a commitment to make floating euro payments. In entering into this swap, BNP takes on mortality exposure, which it then hedges with Partner Re. Thus, if Partner Re defaults, that is BNP's problem, and if BNP defaults, that is the EIB's problem. However, EIB is still committed to pay investors regardless of whether Partner Re or BNP default or not.

4.3.9 For their part, investors have the protection of the EIB's commitment to repay, backed by the EIB's AAA credit rating. For its part, the EIB has the protection of BNP's commitment to take on the bond's longevity risk exposure, and this commitment is backed by BNP's AA credit rating and by the knowledge that BNP has reinsured that risk with Partner

Re. For its part, BNP has the protection of the reinsurance provided by Partner Re, whose rating is also AA.

4.3.10 The EIB/BNP longevity bond has some attractive features:

- (i) Its cash flows are designed to help pension plans hedge their exposure to longevity risk. To be more precise, they provide an ideal hedge against a notional annuity provider who is committed to providing level annuity payments to the reference population over a horizon of 25 years.
- (ii) The survivor index $S(t)$ is calculated with reference to crude death rates published by the ONS. These death rates are a reliable and easily obtainable public source. This helps reassure investors that they would have full access to the data and would not lose out as a result of insurance companies manipulating their reported death rates. The use of crude death rates also avoids arguments over smoothing methodologies.
- (iii) Trends in national mortality should provide a reasonable match for trends in annuitants' mortality, and thus help to reduce basis risk in an annuity book that might be hedged by an investment in the longevity bond.

4.3.11 As noted earlier in ¶4.3.1, the EIB/BNP longevity bond was only partially subscribed and was later withdrawn for redesign. There seem to be a number of reasons for its slow take up and perhaps lessons can be learned for future contract design:

- (i) It is likely that a bond with a 25 year horizon provides a less effective hedge than a bond with a longer horizon. (Evidence to this effect is provided by Dowd, Cairns & Blake, 2005.) Similarly, the bond might prove to be a less effective hedge for pension liabilities linked to different age cohorts or to females. This means that the EIB bond might not be a particularly effective hedge for the kind of annuity book for which it was designed, and this consideration might have discouraged annuity providers from investing in it.
- (ii) The amount of capital required is high relative to the reduction in risk exposure. This makes the BNP bond capital-expensive as a risk management tool.
- (iii) The degree of model and parameter risk is quite high for a bond of this duration (see, for example, Cairns, Blake & Dowd, 2005a), and this degree of uncertainty might make potential investors and issuers uncomfortable. Thus, even if the bond provides a perfect hedge, there will be uncertainty over what the right price to pay or charge should be.
- (iv) Potential hedgers might feel that the level of basis risk is too high relative to the price being charged. For example, basis risk can arise because annuitants are likely to experience more rapid mortality

improvements than is reflected in the population-wide index on which the payments are determined. Basis risk can also arise because the longevity bond specifies level annuity payments, whereas most real-world pension schemes allow for escalating (i.e. inflation-linked) payments. A further cause for basis risk is inaccuracy in the estimates of number of deaths (e.g. people dying while on holiday, slow notification of pensioner death) or in the number exposed to risk (e.g. the number exposed to risk is based on population projections from the last census date), or a failure to ensure these correspond.

- (v) The underlying index is calculated with reference to central death rates. However, the use of central death rates means that $S(t)$ will underestimate the true proportion of the cohort that survive. A more natural definition for the survivor index, which avoids this bias, would make reference to mortality rates: that is, $S(t) = S(0) \times (1 - q(2003, 65)) \times (1 - q(2004, 66)) \times \dots \times (1 - q(2002 + t, 64 + t))$ where the $q(y, x)$ are mortality rates for age x in year y .

4.3.12 Problem (i) can be addressed by increasing the maturity of the longevity bond. One objection sometimes made to a longer maturity longevity bond was that gilts were themselves limited to 25 year maximum maturities, and the absence of ‘ultra-long’ gilts made it difficult for financial institutions to deal in ultra-long bonds of any sort themselves (because of market illiquidity, hedging problems, etc.). This was a reasonable argument, but a recent change in DMO policy makes this argument harder to sustain: in March 2005, the DMO announced that it will start to issue ‘ultra-long’ gilts, and as the volume of these bonds increases, it will become easier for financial institutions to obtain safe ultra-long bonds that they can use for financial engineering or hedging purposes. We might therefore expect to see longer maturity longevity bonds and expect these to be better received. (However, there is also the secondary consideration that a maturity of 25 years might also have been chosen because individual-age mortality rates are available only up to age 89, making longer terms more difficult to handle.) We have more to say on very long maturity longevity bonds in Section 6.1.

4.3.13 Problem (ii) is far from unique and was a notable problem with many non-life securitisations when they were first issued. The answer is to find ways of increasing gearing to provide the same exposure to risk for a lower initial capital outlay. The experience with non-life insurance securitisations is also reassuring in this regard, as they too saw a shift towards more highly geared contracts as the market developed. Several structures that gear up the exposure to longevity risk follow in Sections 6.4 to 6.5.

4.3.14 Problem (iii) has also arisen many times before. For example, it arose when index-linked gilts were first issued in the U.K. in the early 1980s, when investors had little real idea of the data-generation process underlying the RPI inflation rate. However, this did not stop index-linked gilts from

becoming a very successful innovation, and over a fifth of gilts outstanding are now index-linked. In the current context, therefore, after a small number of longevity bonds have been issued investors will be more comfortable with how to price further issues and judge the value of existing securities.

4.3.15 Problem (iv) is discussed further in Section 10.7.

4.3.16 Problem (v) can be handled in a number of different ways. For example, if one wishes to avoid the introduction of subjectivity into how the $q(y, x)$ are calculated, one could specify the usual approximation $q(y, x) = m(y, x)/(1 + \frac{1}{2}m(y, x))$ within the terms of the contract. Alternative methods for inferring the $q(y, x)$ from the crude death rates could be specified. However, it is convenient if we defer any longer discussion of index issues to Section 10 below.

5. NEW MORTALITY-LINKED SECURITIES

5.1 We will now describe some new mortality-linked securities. Broadly speaking, these can be classified into various types of: (i) longevity bond, (ii) mortality swap, (iii) mortality futures, and (iv) mortality options.

5.2 These securities have the usual features we would expect of bonds, swaps, futures, and options. In particular, there is the distinction between those that are traded over-the-counter (e.g. swaps), and those that are traded in organised exchanges (e.g. futures). The former have the attraction that they can be tailor-made to the requirements of a user (which keeps down basis risk), but have the disadvantage that they have thin secondary markets (which makes positions harder to unwind); the latter have the attraction of greater market liquidity (which facilitates unwinding), but have the disadvantage of greater basis risk.

5.3 There are also the usual credit risk issues to consider. With exchange-traded securities, credit risk is handled by the exchange itself standing between traders as the opposite counterparty to each transaction: this means that the exchange itself guarantees all trades, and then protects itself by means of margin requirements and other conditions imposed on traders. These arrangements ensure that traders no longer have to worry about each other's credit-worthiness, although this protection comes at the cost of the margin payments and other requirements imposed by the exchange. The situation with OTC securities is very different: these are essentially bilateral deals, and (depending on the type of security) at least one, and possibly both counterparties, acquires a potential credit risk exposure. Credit issues are then potentially very significant, and a whole range of possible arrangements can be made to deal with these issues. We shall have more to say on credit risk mitigation in Section 11.

5.4 We now consider each of these new types of mortality security in turn.

6. LONGEVITY BONDS

6.1 *Categories of Longevity Bond*

There are many possible types of longevity bond, but they fall under two broad categories. The first are ‘principal-at-risk’ longevity bonds, of which the Swiss Re bond is an example. These are longevity bonds in which the investor risks losing all or part of the principal if the relevant mortality event occurs. The second are ‘coupon-based’ longevity bonds, of which the EIB/BNP bond is an example. These are bonds in which the coupon-payment is mortality-dependent. The nature of this dependence can also vary: the payment might be a smooth function of a mortality index, or it might be specified in ‘at risk’ terms, i.e. the investor loses some or all of the coupon if the mortality index crosses some threshold. Since these are designed as hedge instruments, it makes sense that these bonds take the form of annuity bonds, and have no terminal repayment of principal. However, one can also imagine various types of ‘hybrid’ longevity bonds, such as repayment-of-principal longevity bonds in which both principal and coupon are at risk if specified mortality events occur.

6.2 *Classical Longevity Bonds*

There are many possible types of coupon-based longevity bond. A natural one is a ‘classical’ longevity bond of the type first proposed by Blake & Burrows (2001) and given the name survivor bond. This is a longevity bond whose coupon payments are proportional to the survivorship rate of the specified reference population and whose final payments finish, not after 25 years, but at the death of the last surviving member of the reference cohort. So, for example, if the reference cohort is initially aged 65, and if the longest lived member of it survives to an age of 115, then the last payment on the survivor bond would occur after 50 years. A classical longevity bond can also be regarded as having a stochastic maturity, with the stochastic variable being the lifetime of the longest-lived member of the annuitant cohort. Such a bond has the attraction that it provides a better hedge than an EIB/BNP-type bond whose maturity is limited to 25 years.

6.3 *Zero-Coupon Longevity Bonds*

The longevity bonds described above provide a series of annual payments. However, as happened in the gilts market, one can envisage that single-coupon longevity bonds (‘longevity zeros’) might be issued or financially engineered by stripping ‘standard’ longevity bonds. The attraction of zeros is that they provide building blocks for tailor-made positions. A two-dimensional spectrum of such bonds could be issued: one dimension relating to the cohort being followed and the other relating to the maturity date. The availability of a sufficient variety of bonds from this two-dimensional spectrum would then enable insurance companies to construct portfolios of

longevity bonds that provide close fits to the size/age features of their particular annuity books. However, it seems likely that the market for longevity zeros would be quite illiquid as most such bonds would be bought on a buy-and-hold basis.

6.4 Geared Longevity Bonds and Longevity Spreads

6.4.1 We might also have geared longevity bonds which enable users to meet their hedging demands for a much reduced capital outlay.

6.4.2 One way to construct such bonds would be as follows. Looking ahead from time 0, the payment on each date t can in theory range from 0 to 1 (times the initial coupon). However, again looking ahead from time 0, we can also suppose that the payment at time t is likely to fall within a much narrower band, say $S(t) \in [S_l(t), S_u(t)]$. For example, if we are using a stochastic mortality model we could let $S_l(t)$ and $S_u(t)$ be the 2.5% and 97.5% percentiles of the simulated distribution of $S(t)$. These simulated confidence limits become part of the contract specification at time 0.

6.4.3 We now set up a special purpose vehicle (SPV) at time 0 (similar to the arrangement in Figure 4) that holds $S_u(t) - S_l(t)$ units of the fixed interest zero-coupon bond that matures at time t for each $t = 1, \dots, T$ (or its equivalent using floating-rate debt and an interest-rate swap). Suppose the SPV is financed by two investors A and B . At time t , the SPV pays:

- (i) $S(t) - S_l(t)$ to A with a minimum of 0 and a maximum of $S_u(t) - S_l(t)$; and
- (ii) $S_u(t) - S(t)$ to B with a minimum of 0 and a maximum of $S_u(t) - S_l(t)$.

6.4.4 The minimum and maximum payouts at each time to A and B ensure that the payments are always non-negative and can be financed entirely from the proceeds of the fixed-interest zero-coupon bond holdings of the SPV. (The resulting (minor) optionality is reminiscent of the construction in Lin & Cox, 2005.)

6.4.5 The payoff at t can equivalently be written as $(S(t) - S_l(t)) + \max\{S_l(t) - S(t), 0\} - \max\{S(t) - S_u(t), 0\}$: that is, a combination of a long forward contract, a long put option on $S(t)$ (or a ‘floorlet’), and a short call on $S(t)$ (or a ‘caplet’). The bond as a whole, therefore, is a combination of forwards, floorlets and caplets. Continuing with the option terminology, we can also observe that the payoff to investor A is often referred to as a *bull spread*, and for this reason we refer to the payoff in the current context as a *longevity bull spread*. (For a further discussion of *survivor* — or *longevity* — caps and floors, see Section 9.2.)

6.4.6 Let us suppose that, for each t , $S_l(t)$ and $S_u(t)$ have been chosen so that the value of the floorlet and the caplet are equal. In this case, the price payable at time 0 by investor A is equal to the sum of the prices of the T forward contracts paying $S(t) - S_l(t)$ at times $t = 1, \dots, T$. This is equal to (a) the price for the longevity bond paying $S(t)$ at times $t = 1, \dots, T$, minus

(b) the price for the fixed interest bond paying $S_i(t)$ at times $t = 1, \dots, T$. This structure therefore gives investors a similar exposure to the *risks* in $S(t)$ for a lower initial price. For this reason we describe the collection of longevity bull spreads as a geared longevity bond.

6.4.7 As an alternative $S_u(t)$ might be set to 1, meaning that the caplet has zero value. With this structure, investor A has full protection against unanticipated improvements in longevity, but gives away any benefits from poorer longevity than anticipated.

6.4.8 It is important to note in the above construction that there is a smooth progression in the division of the coupon payments between the counterparties over the range of $S(t)$. This is preferable to a contract that has a jump in the amount of the payment as $S(t)$ crosses some threshold. In such contracts as barrier options, arguments can often arise as to whether the particular threshold was crossed or not. Such difficulties are avoided with the smooth progression.

6.5 *Deferred Longevity Bonds*

6.5.1 Another way of increasing gearing is by issuing bonds with deferred payment dates. We noted already that a criticism of the EIB longevity bond is that the early coupon payments have very low longevity risk attached to them, and estimates from Cairns, Blake & Dowd, 2005a, suggest that the first 10 years' cash flows are very low risk. Yet these cash flows are also the most expensive part of the bond. For users wishing to use these bonds as hedging instruments, such bonds use up a lot of capital to cover a long period of low-risk payments. A natural way to deal with this problem is for users to buy longevity bonds with deferred payments. The deferments would save a large amount of capital, and so increase the gearing. This, in turn, would make such longevity bonds much more attractive as hedging instruments.

6.5.2 These deferred longevity bonds can also be regarded as a form of mortality forward contract. As with conventional forwards, one can envisage that they might take a large number of different forms.

6.6 *Principal-at-Risk Longevity Bonds*

This type of bond has a similar structure to the Swiss Re mortality bond. The bond is issued by a single pension plan or annuity provider (A) using a special purpose vehicle (SPV — see Section 11.5). At the outset the SPV is funded by contributions from A and external investors (B). The total outgo of the SPV would mimic either a floating-rate or a fixed interest bond paying annual coupons and with a final repayment of principal at maturity. Under 'normal' circumstances coupons and principal would be payable in full to B . However, if a designated survivor index, $S(t)$, exceeds a specified threshold then a reduction in the repayment of principal to B (and possibly also the coupons) would be triggered, with the residual payable to A . The result of

this is that A benefits financially if longevity improves by more than expected. Alternatively a reduction in the repayment of principal could be linked to some form of weighted average of the $S(t)$ s (a form of Asian option).

7. MORTALITY SWAPS

7.1 *Introduction*

7.1.1 A mortality swap is an agreement to exchange one or more cash flows in the future based on the outcome of at least one (random) survivor or mortality index. We have already met one form of mortality swap in Section 4.3: this was the mortality swap embedded in the EIB bond.

7.1.2 Mortality swaps as defined in the previous paragraph bear considerable similarity to reinsurance contracts. Both often involve ‘swaps’ of anticipated for actual payments (or claims), and both might be used for similar purposes. However, there are major differences between them. Most especially, mortality swaps are not insurance contracts in the legal sense of the term and therefore not affected by some of the distinctive legal features of insurance contracts (e.g. indemnity, insurable interest, etc.). Instead, mortality swaps are subject to the (generally less restrictive) requirements of securities law. So, for example, a mortality swap allows one to speculate on a random variable, whereas an insurance contract does not. Similarly, an insurance contract requires the policyholder to have an insurable interest, but a mortality swap does not.

7.1.3 Mortality swaps could take many different forms and are discussed in some detail by Cox & Lin (2004), Lin & Cox (2005) and by Dowd, Blake, Cairns & Dawson (2005).

7.2 *Attractions of Mortality Swaps*

Mortality swaps have certain advantages over longevity bonds. They can be arranged at lower transactions cost than a bond issue and are more easily cancelled. They are more flexible and they can be tailor-made to suit diverse circumstances. They do not require the existence of a liquid market, just the willingness of counterparties to exploit their comparative advantages or trade views on the development of mortality over time. Mortality swaps also have advantages over traditional insurance arrangements: they involve lower transactions costs and are more flexible than reinsurance treaties, and so on.

7.3 *A Nascent Market in Mortality Swaps*

We know from our industry contacts that some insurance companies have already entered into mortality swaps on an over-the-counter (OTC) basis. The market is in its early stages and concrete details are hard to pin down,

but off-the-record discussions with practitioners indicate that a number of reinsurers are transacting OTC vanilla mortality swaps in which the preset-rate leg is linked to a published mortality projection, and the floating leg is linked to the counterparty's realised mortality. There are also related derivatives being traded that involve the securitisation of life offices' annuity books. Typically, the reinsurers also act in syndicates to spread their exposures. As far as we can tell, the counterparties are usually life companies, but some investment banks are also interested. The attractions of these arrangements are the obvious ones of risk mitigation and capital release for those laying off longevity risk, and low-beta risk exposures for those taking it on.

7.4 *One-Payment Mortality Swaps*

In the most basic case, a mortality swap would involve the exchange of a single preset payment for a single random mortality-dependent payment. More precisely, suppose that at time 0, two firms enter into an agreement to swap a preset amount $K(t)$ for a random amount $S(t)$ at some future time t . As with a conventional forward rate agreement (FRA), $K(t)$ can be interpreted as a coupon associated with an implicit notional principal, and to keep mutual credit risks down, it makes sense for the agreement to specify that the two parties exchange only the net difference between the two payment amounts: so firm A pays firm B an amount $K(t) - S(t)$ if $K(t) > S(t)$ and B pays A an amount $S(t) - K(t)$ if $S(t) > K(t)$. $S(t)$ is related to the number of people from a specified reference population (e.g. the whole population or the number of annuity holders at time 0) who have actually survived to time t . Ex post, A benefits if $S(t)$ turns out to be high relative to $K(t)$ and loses if $S(t)$ turns out to be low: firm A has a long exposure to $S(t)$, whilst B has a short exposure to $S(t)$.

7.5 *Vanilla Mortality Swaps*

We can regard this basic one-payment swap as the core building block in a vanilla mortality swap (VMS), in which the parties agree to swap a series of payments periodically (that is, for every $t = 1, 2, \dots, T$) until the swap matures in period T . A VMS is analogous to a vanilla interest-rate swap (IRS), which involves one fixed leg and one floating leg typically related to a market rate such as LIBOR. However, there are several key differences. The fixed leg of the IRS specifies payments that are constant over time, whereas the corresponding leg of the VMS involves preset payments that decline over time in line with the survivor index anticipated at time 0. Also, the floating leg of the IRS is tied to a market interest rate, whereas the floating leg of the VMS depends on the realised value of the survivor index at time t . Finally, the IRS can be valued using a zero-arbitrage condition because of the existence of a liquid bond market. This is not the case with a VMS which must be valued in an incomplete markets setting.

7.6 *Other Mortality Swaps*

We can also envisage many other types of mortality swap. For example, we could have swaps that involve the exchange of one floating-rate payment for another. We could also imagine more elaborate types of swap: swaps on mortality spreads, cross-currency mortality swaps, mortality swaps in which one or more floating payment depends on a non-mortality random variable (e.g. an interest rate, a stock index, etc.), and mortality swaps with embedded features such as options.

7.7 *Uses of Mortality Swaps*

7.7.1 Mortality swaps have a number of possible uses. One insurer, wishing to manage the risks on its annuity book, might be on one side of the swap, while on the other side of the swap might be a capital market institution wishing to acquire longevity risk exposure. Swaps are attractive because even where alternatives exist, swaps often offer the parties concerned less costly ways of managing their longevity risks, which can also be tailor-made to the users' requirements. Should circumstances change and one party later wish to change its desired exposure, a swap also gives it a much more flexible means of altering its exposure.

7.7.2 As Cox & Lin (2004) explain, a mortality swap can also be used to help firms that run both annuity and life books manage the natural hedges implicit in their positions. The type of swap in this case might be a floating-for-floating swap, with one floating leg tied to the annuity provider's annuity payments and the other to the life assurer's insurance payouts.

7.7.3 There are also other ways in which mortality swaps can be used to manage mortality exposures. Bearing in mind that swap payments would be conditioned on particular time periods and reference populations, firms could use such swaps to manage their exposures across both reference populations and across the 'mortality term structure'. For example, firms in different countries could enter into such swaps to diversify their longevity risk exposures; alternatively, firms might enter into such swaps to alter their 'mortality term' risk exposures.

7.7.4 Finally, mortality swaps also have their uses as vehicles to speculate on longevity risk.

8. MORTALITY FUTURES

8.1 *Introduction*

8.1.1 We have seen above a range of traded securities whose cash flows are linked to the development of mortality over time. In this section and the next we will consider the special class of securities that might be traded in a futures and options exchange. We consider first futures contracts.

8.1.2 The basic form of a futures contract involves defining (a) the

underlying (typically price) process, $X(t)$, that will define the payoff on the futures contract, and (b) the delivery date, T , of the contract. Sometimes elements of optionality will be incorporated into a futures contract, such as a range of delivery dates instead of a single date, but these tend to be relatively minor in nature and in value. For notational simplicity here we will define one trading day as the time unit. The mechanics of a futures market work as follows:

- (i) Let $F(t, T)$ be the futures price at t for delivery of $X(T)$ at time T .
- (ii) By definition, a futures contract has zero value at each time t . However, to minimise credit risk both counterparties are required to place collateral into a margin account that will absorb normal movements in prices.
- (iii) At time $t + 1$ the revised futures price will be $F(t + 1, T)$. This change in price requires a margin payment of $F(t + 1, T) - F(t, T)$ to the holder of the long position in the futures contract. The value of the contract is then reset to zero. If necessary one or other margin account will need to be topped up, or the contract closed out.
- (iv) At time T , the final margin payment of $F(T, T) - F(T - 1, T) = X(T) - F(T - 1, T)$ is equivalent in cash terms to delivery of the underlying $X(T)$ in return for the price $F(T - 1, T)$.

8.1.3 The challenge in the current context is to consider whether or not there exists one or more suitable underlying variables, $X(t)$, that will result in a successful futures market. Where $X(t)$ itself represents the price of a traded asset, the advantage of a futures market is normally that it allows stakeholders to trade in the underlying risk with lower transaction costs and in a market with greater liquidity than is usually possible from trading in the underlying the spot market.

8.1.4 We will consider first the factors that contribute to a successful market or otherwise and then look at the existing annuities market and a possible market in longevity bonds as providing us with possible underlying processes.

8.2 *Factors Making for the Success of Futures Contracts*

A large number of studies — see, for example, Gray (1978), Ederington (1979), Carlton (1984), Black (1986), Pierog & Stein (1989), Corkish, Holland & Vila (1997) and Brorsen & Fofana (2001) — suggest that the following factors are key to the success of a futures contract, where a successful contract is defined as one that has a consistently high daily volume of trade and open interest (the number of contracts carried overnight, i.e. not closed out before close of business):

- (i) There must be a large, active and liquid spot market for the underlying with good price transparency. This is by far the most important factor: indeed it is extremely rare for a futures contract to survive without a

spot market satisfying these conditions. One example of a futures contract that has survived with no traded underlying is weather futures. The contract nevertheless suffers from serious problems, the most important of which is location basis risk, arising from the different weather experiences of the index station and the hedged location (Rohrer, 2004).

- (ii) Spot prices (i.e. the traded prices of the underlying security) must be sufficiently volatile to create both hedging needs and speculative interest. Bessembinder & Seguin (1993) argue that relative hedging demand can be measured by the level of open interest relative to volume, since the former excludes the many speculators who do not hold overnight positions. A low open interest to volume ratio is an indication of high liquidity which is another sign of a successful futures contract.
- (iii) The underlying must be homogeneous or have a well-defined grading system. Homogeneity implies that different units of the underlying are identical, so that, if the contract goes to delivery, the buyer knows precisely what he will receive. This condition does not generally apply to commodity contracts (such as wheat or oil futures). Different batches of the underlying will have different grades or quality. However, these are standard problems with many futures contracts (e.g. commodity futures) and can be resolved if batches of the underlying can be assigned to an agreed point on an acceptable grading system.
- (iv) The market in either the underlying or the futures must not be heavily concentrated on either the buy or sell side, because this can lead to price manipulation. An example of a market that is known for this is the silver futures market. A famous attempt to corner the market and artificially raise prices was made by the Hunt brothers in 1979-80.
- (v) Liquidity costs (i.e. bid-ask spreads and execution risk — the risk of adverse price movements before trade execution) in the futures contract must not be significantly higher than those operating in any existing cross-hedge futures contract.
- (vi) A successful futures market requires the active participation of both hedgers and speculators. Hedgers require that the futures contract provide them with an effective tool for risk reduction. Speculators require sufficient liquidity and volatility in the market to make short-term trading possible.

8.3 *Some Instructive Failures*

8.3.1 It is also instructive to examine why some specific futures contracts failed or never started.

8.3.2 A famous example was the U.S. CPI (consumer price index) futures contract listed on the U.S. Coffee, Sugar and Cocoa Exchange in June 1985. This contract was delisted in April 1987 with only 10,000 contracts ever having been traded. (A successful financial futures contract

could expect to have a *daily* volume of 10,000 contracts!) Srinivasan (2004) suggests that the failure of this contract was due to a combination of the absence (in the U.S.A.) of an inflation-linked securities market at the time, the fact that the underlying was an infrequently published index, and the fact that there was no stable pricing relationship with other instruments. A second CPI futures contract to fail was the Treasury inflation-protected securities (TIPS) futures contract listed on the Chicago Board of Trade in June 1997 and subsequently delisted before the end of the year with only 22 contracts ever traded. Srinivasan (2004) explains this failure as partly due to the contract being premature, as the underlying TIPS had only started trading five months before and there was only a single 10-year TIPS outstanding. However, he also suggests that its failure might also have been due to the fact that the futures contract competed with the underlying for liquidity, and to uncertainty over the fate of the TIPS programme overall.

8.3.3 The validity of these conclusions would also appear to be confirmed by the apparent success of a CME CPI futures contract launched in February 2004. Srinivasan (2004) suggests that the success of this contract was due in part to the fact that inflation-linked securities have gained acceptance amongst investors and TIPS have evolved into a recognised asset class. In addition, the U.S. Treasury is committed to long-term TIPS issuance, CPI futures use the same inflation index as TIPS, complement TIPS rather than compete with it, and the contract is traded on the Globex electronic trading platform and has automated two-sided price quotes from a leading market maker. These factors are also helpful because they allow for well understood pricing relationships between TIPS, fixed-interest Treasury bonds and CPI futures.

8.3.4 Another instructive example is the set of four property futures contracts introduced in May 1991 by the London Futures and Options Exchange (FOX). These contracts comprised a residential property contract, a mortgage interest rate (MIR) contract, a commercial property capital value contract and a commercial property rent contract. The indices for the underlying were the Nationwide Anglia Building Society house price (NAHP) index, the FOX MIR index, the Investment Property Databank (IPD) capital value index and the IPD commercial rent index, respectively. However, trading was extremely thin and all the contracts were suspended in October 1991. Patel (1994) reports that none of the contracts reached economically viable trading volume. There appear to have been two main reasons for the failures of these contracts. First, there were problems with the construction of the indices. In the case of the property and rental contracts, there was lag dependence over time caused by property price and rental changes spilling over gradually from region to region, taking up to a year to ripple completely across the country, starting from London and the South East. This meant that the NAHP index always lagged behind the market. In the case of the MIR contract, there were time delays in processing new

mortgage loans. These factors, in turn, would have created asymmetric information (i.e. led to some parties being better informed than others) which would also have discouraged participation in the FOX markets. Second, there was time basis risk. Given the illiquid nature of the spot market and the long time lags and high transactions costs involved in buying and selling houses, it was hard to determine which maturity of futures contract should be used for hedging purposes. Further, arbitrageurs were unable to correct for pricing anomalies because of restrictions on short-selling a portfolio of property. These factors led to the indices being false indicators of true market prices and to anomalous differences between spot and futures prices persisting.

8.4 Annuity Futures

8.4.1 The concept of an annuity futures market was introduced briefly by Cairns, Blake & Dowd (2004) and we now discuss this in more detail. The basic idea of an annuity futures contract is a simple one. Suppose that $AR_J(t, x, y, p)$ represents the market joint-life level annuity rate at time t per £50,000 single premium for a female non-smoker aged x and a male non-smoker aged y (that is, the single premium of £50,000 will purchase an annuity of $AR_J(t, x, y, p)$ per annum while both are alive, reducing to $p \cdot AR_J(t, x, y, p)$ per annum on the first death; typical values for p being $\frac{2}{3}$ and $\frac{1}{2}$). The equivalent notation for single-life level annuities will be denoted by $AR_M(t, x)$ and $AR_F(t, x)$ for males and females respectively. For convenience we will abbreviate this cumbersome notation to $AR(t)$ for a specific type of annuity (e.g. a single-life, level annuity to a male non-smoker aged 65). A futures contract would have $AR(t)$ as its underlying index and there would be a variety of maturity dates.

8.4.2 However, there are a variety of substantive issues that would need to be resolved. These relate to the illiquidity and inefficiency of the spot market for annuities. The market for immediate annuities in the U.K. is a fairly active one, but it is also both illiquid and inefficient. Individual insurers change their market annuity rates relatively infrequently, instead of changing their rates on, for example, a daily basis to reflect changes in the gilts yield curve. So even if insurer A is quoting the best price in the market there will be times when a prospective annuitant might be advised to delay purchase of an annuity because it is expected that A's annuity rate will soon move in a favourable direction to reflect changes to the gilts yield curve that have already happened. In addition, a large proportion of the buyers of annuities (the individual policyholders) do not make the effort to seek out the best rates. Often a personal pension plan member will purchase an annuity from the same insurer with whom they have held their pension savings prior to retirement. We infer from this (without further evidence) that insurers who quote uncompetitive annuity rates do still get business from such policyholders. Related to this, insurers who are active players in the market

can be seen to move in and out of the market. For example, an insurer might move to less competitive prices if, in a given year, they have reached their target for new annuity business. Such targets might exist within insurers who aim to keep a balance between annuity and term-assurance business.

8.4.3 All of these points can be seen in Figure 5. In this figure we have plotted the market joint-life annuity rates for a male and female aged 65 — $AR_j(t, 65, 65, \frac{2}{3})$ — for 9 insurers over the period July 2002 to November 2004. The best annuity rate is compared with gilts yields in Figure 6. We can make the following observations:

- (i) Company A can be seen to be an active player in the market. A is often quoting the best rate but not always.
- (ii) Company C is also an active player but seems to move in and out of the market to a greater extent. For example, around July 2004 C's rates dropped significantly at a time when no other company's did.
- (iii) Company E can be seen to have been a fairly active player in 2002 and the early part of 2003. However, after that, their business strategy clearly changed and they have quoted a very low rate since the middle of 2003.
- (iv) In the middle of the market we can see companies F and I. In both cases one might assume that the majority of their annuity business comes from existing policyholders rather than new policyholders on the open market. F's strategy appears to be to keep its price stable and fairly uncompetitive. In contrast, I is also uncompetitive but nevertheless changes its annuity rates in line with, but consistently below, the active players in the market.
- (v) The market has become more heterogeneous since 2002. In July 2002 the difference between the top and the bottom rates was about 10%. In the latter part of 2003 and 2004 this gap had widened to 20%.

8.4.4 In Figure 6 we can see how market annuity rates do to a large extent track changes in interest rates.

- (i) The short-term increases in interest rates at A and B were not reflected in annuity rates, whereas after the middle of 2003 the accuracy of tracking was higher with the peaks in interest rates at C and D being reflected in market annuity rates.
- (ii) From the middle of 2003 to the end of 2004 the correlation between yields and annuity rates was 0.72 (0.80 if a time lag of 2 weeks is used).
- (iii) Fluctuations since the middle of 2003 are consistent with the joint-life annuity having a duration of about 9 years.
- (iv) The steep drop in annuity rates in 2002 is partly due to falling interest rates. However, a comparison with fluctuations in 2003 and 2004 suggests that a large part of the fall in 2002 must be due to other factors such as revisions to the mortality tables being used by insurers.

8.4.5 One could argue that the current relative inefficiency of the

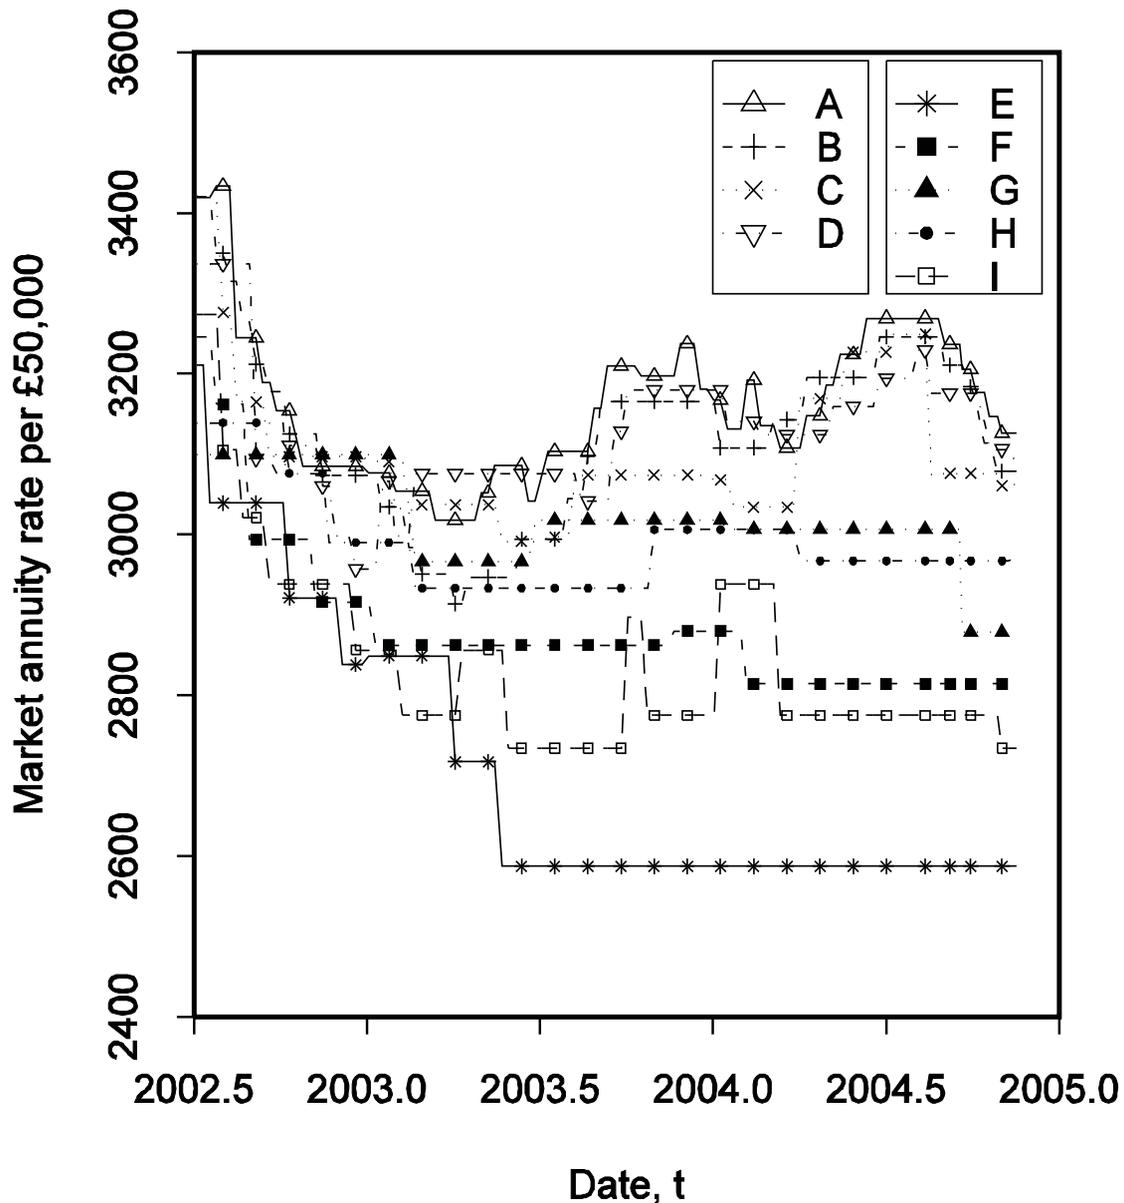


Figure 5. Market annuity rates, $AR_J(t, 65, 65, \frac{2}{3})$, available for a single premium of £50,000 from nine U.K. insurers from the middle of 2002 ($t = 2002.5$) to November 2004, level, joint-life annuity payable to male and female lives both aged 65 for 5 years guaranteed, reducing to two-thirds on the first death. Source: data compiled by Aviva and provided through AFPen

annuities market might make an annuity futures market attractive to speculators. Specifically if they can devise a means of exploiting the slow reactions of insurers to changes in interest rates then there will be profits to be made.

8.4.6 All of this inefficiency in the immediate annuities market means that defining a suitable index underlying the futures contract is difficult.

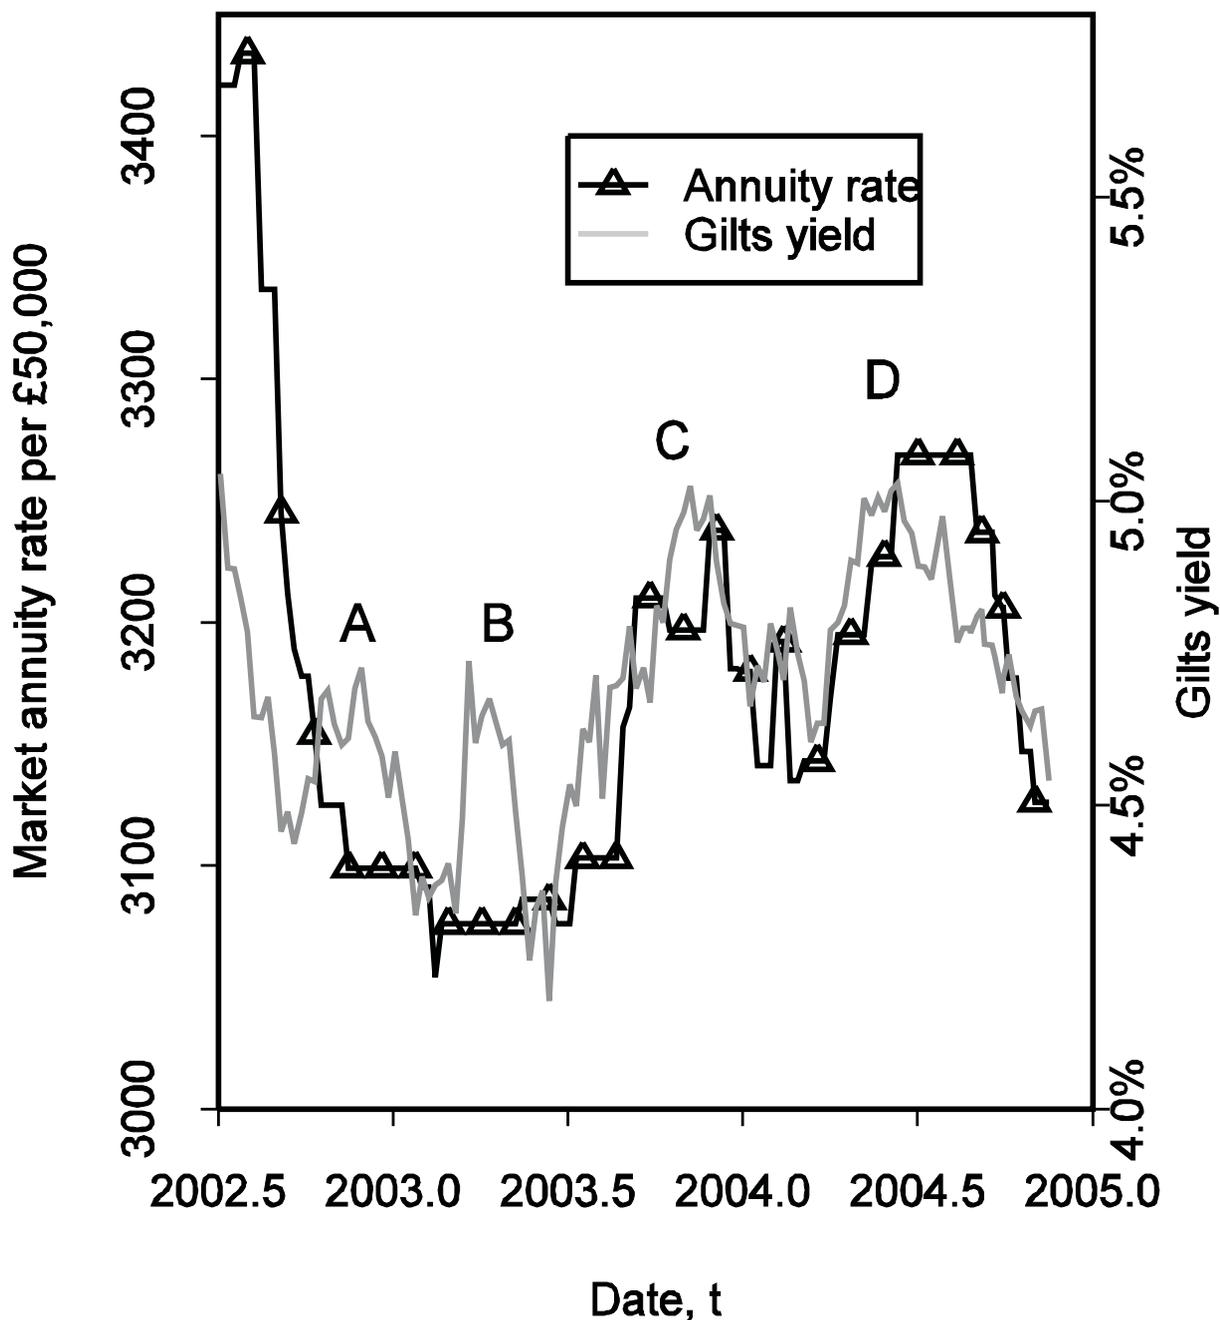


Figure 6. Comparison of the best market annuity rate (left-hand axis) with the gilts yield (right-hand axis). Source: data compiled by Aviva and provided through AFPen

Clearly the index needs to be derived on a daily basis from the market annuity rates of all of the insurers. However, a number of questions need to be asked:

- (i) What weight should be placed on each insurer's rate? Greater weight should be given to the most competitive prices, but should the medium and low-ranked insurers be given a low weight or zero weight?

- (ii) Should these weights change over time? The answer is clearly yes: if an insurer decides to move out of the market then its weight should change to reflect this. However, should the weights vary in a way that minimises the impact of an individual company deciding to move in or out of the market?
- (iii) For an annuity futures contract to be successful, it is important to have a clear picture of who the hedgers are. Is it individual pension plan members? Is it large defined contribution pension plans? Is it large final salary pension plans? The underlying annuity index then must be designed to cater most accurately for the hedging requirements of this specific group.
- (iv) Furthermore, in answering these questions, we would also have to take account of a possible objection: since annuity futures would be primarily driven by interest rate changes, and since these can be hedged using traditional fixed-income instruments, one has to ask why anyone but individual pension-plan members would ever want to use these futures for hedging purposes.
- (v) There is also another problem. There is no obligation on insurers to reveal their prices, and it may not be in their interest to do so. And, if insurers get better in the future at segmenting markets and/or avoiding price disclosure (as seems plausible), then an annuity futures based on available data might be quite uninformative.

8.5 *Longevity Futures*

8.5.1 *Longevity bond prices as the underlying*

8.5.1.1 If a liquid market in longevity bonds develops in time then it might be possible for a futures market to develop which uses the price or prices of longevity bonds as the underlying. The nearest equivalent to this in the U.K. is the LIFFE long-gilts futures contract.

8.5.1.2 A key issue to consider here is the likely involvement of speculators and arbitrageurs. Day-to-day volatility in longevity bond prices will be driven by changes in interest rates, whereas the risk associated with changes in longevity emerges over longer periods of time. It follows that speculators and arbitrageurs will have to feel that interest-rate movements in the gilts market are not being properly reflected in the longevity bonds market in order for them to participate in the new futures market.

8.5.2 *Survivor indices as the underlying*

8.5.2.1 The previous section used the prices of traded longevity bonds as the underlying. If we look outside this constraint then a possible approach is to use a survivor index as the underlying, one example of this being the index that underpins cash flows in the EIB longevity bond. This type of contract mimics the CME CPI futures market. There is no market in the CPI

itself but there is a liquid TIPS market. It was noted above that the TIPS market is central to the success of the new CPI futures market since the combination of the fixed-income and TIPS markets allows partial trading of views on the CPI. Similarly here, if a market develops in longevity bonds, then the combination of this with the fixed-income market might make a futures market in the underlying survivor index attractive.

8.5.2.2 The potential range of indices here is high compared with the CPI contract, with the possibility of one male and one female index for each age. On the one hand, the use of too many indices might reduce liquidity. On the other hand, speculators might wish to trade one futures contract off against another if they take a particular view on how mortality rates at different ages might develop relative to each other.

9. MORTALITY OPTIONS

9.1 *Options*

9.1.1 Options give payoffs that are non-linear functions of underlying variables, and a natural first question with options is why market participants would prefer the non-linear payoffs they generate over the (broadly) linear payoffs of, say, annuity futures. A key part of the answer must be that options might be useful to (a) hedgers who might wish to protect their downside exposure, but leave any upside potential, and (b) speculators who want to trade views on volatility rather than views on the level of mortality (or related, e.g. annuity) rates. For both these purposes, options are (usually) the best type of instrument.

9.1.2 The valuation of options as well as their risk management requires the use of a good stochastic mortality model. A full discussion of this important issue is beyond the scope of the present paper, but as examples the reader is referred to CMI Working Paper 15 (CMI, 2005), Olivier & Jeffery (2004), Cairns, Blake & Dowd (2005a), and references therein.

9.2 *Survivor Caps and Floors*

9.2.1 A possible market in survivor caplets and floorlets has already been alluded to in Section 6.4 on the geared longevity bond. The basic idea is to use a survivor index $S(t, x)$ as the underlying. Now let $s_c(t)$ be the cap rate for exercise date t . The caplet pays $\max\{S(t, x) - s_c(t), 0\}$ at time t . Similarly a floorlet would pay $\max\{s_f(t) - S(t, x), 0\}$. Survivor caplets and floorlets then get packaged into survivor caps and floors. (Alternative names might be *longevity caps and floors*.)

9.2.2 As an alternative to the use of $S(t, x)$, we could use the survivor futures price as the underlying.

9.3 *Annuity Futures Options*

We have described above the annuity futures market. A natural extension to this market is to include options on the annuity futures.

9.4 *OTC Options and Embedded Options*

We described earlier possible OTC contracts, including mortality swaps. It seems likely and indeed desirable that OTC options might also be struck between hedgers and their counterparties. For example, a hedger might arrange for a survivor cap tailored to its own mortality experience (and this arrangement might also be described as a reinsurance deal!).

9.5 *Mortality Swaptions*

9.5.1 A more sophisticated contract would be a mortality swaption. In the case of a mortality swaption, the underlying instrument would be a mortality swap of specified type and maturity. The swaption might be American, European or Bermudan in nature, and would give the holder the right to enter into the swap on one or other side. If the underlying swap is a vanilla mortality swap, the swaption might be a payer swaption, giving the holder the right to enter as the fixed-rate payer, or a receiver swaption, giving the holder the right to enter as the fixed-rate receiver. As with conventional swaptions, a payer swaption can be regarded as a put on survivor rates, because its value would go up when survivor rates fall, and a receiver swaption can be regarded as a call on survivor rates, because its value would increase when survivor rates rise.

9.5.2 With interest-rate swaptions, the intrinsic value of the swap at the exercise date is straightforward to establish with reference to the current yield curve. However, with a mortality swap, the decision is more interesting, since (at least at the current time) there is no liquid market in mortality-linked securities that would allow us to pinpoint the value of the swap at the exercise date.

9.5.3 Mortality swaptions can be used for various risk management purposes, and an obvious use is to provide the option to lock-in future swap rates. An example would be to assist insurance companies in managing the risks of positions in instruments such as guaranteed annuity options (GAOs). Such swaptions would presumably have long terms to maturity, because it is the longer-term longevity risk that, for example, annuity providers are most concerned about.

10. WHICH MORTALITY INDEX?

10.1 *Introduction*

The choice of mortality index is critical to the success or otherwise of a new security, and in this section we will look at this issue in more detail.

10.2 *Frequency of Mortality Data*

10.2.1 One of the first things we must appreciate about mortality data is its lack of timeliness. Whereas financial data is produced on an almost continuous real-time basis (or at worst a daily basis), and economic data, such as price inflation, is typically produced on a monthly basis, mortality data is published much less frequently and is also subject to incurred-but-not-reported (IBNR) issues. At best, it is published only annually, and sometimes it is only published at intervals of 4 or 5 years. Thus, mortality data is anything but timely, and any security linked to a particular mortality table must be designed with this lack of timeliness in mind.

10.2.2 In addition, mortality-linked securities might sometimes rely on unpublished mortality tables that change at irregular intervals. For example, the annuity futures contract depends on the immediate annuity mortality tables in use by several insurers. Each insurer will change its mortality assumptions at different times and these changes will be reflected in changes in the market annuity rates.

10.2.3 On the other hand, annuity rates are published daily and this permits the design of a security that has much greater flexibility with regard to payment dates than securities that are dependent in an explicit way on a mortality table that is published only, say, annually.

10.3 *Choice of Mortality Table*

10.3.1 A number of choices are available for the underlying mortality table.

10.3.2 *National population mortality rates.*

10.3.2.1 For example, the U.K. Office for National Statistics (ONS) publishes tables of mortality for individual calendar years for England and Wales (see www.statistics.gov.uk/StatBase/Product.asp?vlnk=620). The tables are produced annually and are published typically about 18 months after the end of the calendar year of observation. The published tables include crude death rates, rather than smoothed rates. Amongst all mortality tables in the U.K., the ONS table might be considered as the most reliable, since it covers such a large population, and is produced by a trustworthy government department. These statistics are used in the EIB longevity bond.

10.3.2.2 These annual rates of mortality by age can be supplemented by the aggregate mortality data that is available on a weekly basis from the ONS (see www.statistics.gov.uk/StatBase/ssdataset.asp?vlnk=6157). This is a rather crude set of statistics but the relatively high frequency of the releases might encourage shorter-term speculation in mortality-linked securities, thereby increasing liquidity.

10.3.3 *CMI tables*

10.3.3.1 Ungraduated mortality tables are produced by the Continuous

Mortality Investigation Bureau (CMI) generally once every 4 years. Individual tables are produced for different lines of business, e.g. term assurance, immediate annuitants etc., based on mortality returns from the contributing life offices. Experienced mortality is published in 5-year age bands, typically, rather than by individual year of age. The CMI publishes smoothed (or graduated) mortality rates (see Section 10.6 below) at intervals of 8 to 12 years covering the contributing life offices in aggregate. Mortality rates are also analysed at the level of individual companies but these are treated as confidential and not made public. Crude mortality rates are published typically two to three years after the end of the period of observation and smoothed mortality rates some months after that.

10.3.3.2 The CMI tables have the advantage over the ONS tables that the underlying population in many cases is closer in its characteristics to those of the organisations that seek to hedge their exposure to longevity risk, resulting, potentially, in lower basis risk (see Section 10.7 below). However, the CMI tables have a number of drawbacks:

- (i) The underlying population might change in its characteristics over time reflecting the changing preferences of the general population with regard to the provision of insurance and annuities. These changes might not be mimicked in the population that concerns a hedger of longevity risk.
- (ii) The population size in typical CMI studies is typically much smaller than the ONS population statistics. As a result crude mortality rates are less reliable as estimates of the ‘true’ underlying mortality rates. A similar reservation might also apply to smoothed rates.
- (iii) There is a longer time lag between the period of observation and publication of mortality rates.
- (iv) For particular lines of business, one or two insurers can dominate the statistics. Thus, insiders within these companies can observe very much earlier the mortality trends that will ultimately affect contracted payments on a mortality-linked security. Although this is a form of moral hazard, it might not be that substantial given the long timescales involved. However, there might be a perception amongst investors in a mortality-linked security (particularly external, short-term investors) that this moral hazard is significant. Moral hazard is increased further by the fact that individual offices can choose not to submit data to the CMI on certain lines of business. A plausible scenario, therefore, arises where a company could manipulate the CMI table underlying a mortality-linked security by choosing not to submit relevant mortality data.
- (v) The methods used by the CMI to smooth or graduate the crude mortality data change from time to time. This updating ensures that the latest statistical methods and computing power can be applied to produce the next set of tables; however, changes in the calculation methodology introduce uncertainty (or a possible perception of uncertainty), which may put off potential investors.

10.3.4 *The hedger's own mortality experience.*

10.3.4.1 A new security might be linked to the mortality experience of a specific organisation. This will clearly allow this organisation to minimise its own longevity risk, but also creates a high degree of moral hazard that could deter potential investors. For example, Swiss Re could have linked the issue of the Vita Capital mortality bond to their policyholders' mortality experience. However, the bond was intended as a traded security and the best advice they got was to link payments to national mortality indices. Despite this constraint, they were able to tailor the underlying mortality index to match as closely as possible their own exposure to mortality risk through the choice of national market and gender weights, thereby reducing basis risk.

10.3.4.2 In contrast, some OTC contracts do link cashflows to the hedger's own experience: for example, a mortality swap. It is then the responsibility of the two counterparties to ensure that appropriate mechanisms are in place to ensure that the relevant mortality rates are calculated accurately.

10.4 *Information*

For traded securities it is important that potential investors have access to as much relevant information as possible. The less information investors have, or the less forthcoming the issuer appears to be, the more risky the security might seem, and the less likely it is that they will invest. Thus, it is not untypical for offer documents to contain substantial historical data for potential investors to analyse and assess the risks associated with the security as well as its value.

10.5 *Integrity and Moral Hazard*

10.5.1 We have already made some brief remarks on integrity and moral hazard. The index underlying a mortality-linked security needs to have, and must be perceived to have, integrity in the way that it is calculated. If issuers fail to satisfy potential investors on this question, then they run a major risk that the contract will fail to launch.

10.5.2 There are also issues of moral hazard, and moral hazard can take two forms. First, moral hazard exists when, for example, 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. Second, moral hazard exists when there is the possibility for the underlying index to be manipulated (see our comments in Section 10.3.3.2 about the contributors of CMI data). This might exist at the level of crude data provision, or at the point that mortality data is smoothed.

10.6 *Crude or Smoothed Mortality Rates*

10.6.1 Many of the securities discussed above can be linked either to a

table of crude mortality rates or a table of smoothed (or graduated) mortality rates.

10.6.2 Crude mortality rates (or the crude death rates published by most government statistical agencies) give precise statistics about the number of deaths at given ages as a proportion of the number of people alive at that age. A mortality-linked security can make reference to such tables in a straightforward manner. This increases the transparency of the contract and the confidence that investors have in the reliability of the index. However there are some drawbacks.

- (i) Let $\varepsilon(y, x)$ be the number of lives aged x alive at the start of year y , and $q(y, x)$ be the true probability of death during the coming year. Even if $q(y, x)$ is known, the actual number of deaths, $D(y, x)$, is random, and has a binomial distribution with parameters $\varepsilon(y, x)$ and $q(y, x)$. The crude mortality rate is $\hat{q}(y, x) = D(y, x)/\varepsilon(y, x)$ which will hopefully be close to $q(y, x)$, but typically will not be equal. In many circumstances hedgers will be aiming to match the development of $q(y, x)$ over time, or derived quantities such as a survivor index. Differences between $q(y, x)$ and $\hat{q}(y, x)$ will be greater at young ages, where mortality rates are low, and at very high ages, where the number of people alive is relatively small. If a mortality index is calculated with reference to the crude mortality rates rather than ‘true’ mortality rates then an additional source of risk will be added from the perspective of the hedger. In the case of a survivor index these ‘sampling’ errors will be compounded over time.
- (ii) At higher ages, the reported age at death becomes unreliable, leading to larger errors in estimates of mortality rates at those ages.

10.6.3 Smoothing of mortality rates aims to reduce the differences between the true and the crude mortality rates. Examples of smoothing techniques can be found in Currie, Durban & Eilers (2004), and Forfar, McCutcheon & Wilkie (1988). The use of smoothing techniques exploits the fact that adjacent ages (and also calendar years) are subject to very similar rates of mortality. The smoothed mortality rates, $\bar{q}(y, x)$, should be closer to the true values, $q(y, x)$. Potential drawbacks are:

- (i) There are many different techniques that can be used to smooth mortality rates. Over time, new methods are developed that exploit new statistical theory or enhanced computing power. These developments should mean that the smoothed rates should be closer to the true rates than before. These developments are good in some regards, but, in the context of a mortality-linked security, changes in the calculation methodology adds uncertainty and creates doubt in the minds of potential investors over the integrity of the underlying index. As a consequence, the issue documentation of a mortality-linked security should detail precisely how the mortality rates will be smoothed.

- (ii) The use of any smoothing techniques means that the index is less transparent than an index that is calculated directly from crude mortality rates. Less transparency increases the likelihood that an issue will be unsuccessful.
- (iii) Smoothing of mortality rates often takes some time. For example, there can often be a substantial time lag between the publication of crude CMI mortality tables and smoothed CMI tables. The further delay increases basis risk (see Section 10.7) and can increase the likelihood that the issue will be unsuccessful. Of course these delays can be reduced if the issue documentation specifies precisely how the crude death rates are to be smoothed.
- (iv) Crude mortality data sets often have relatively little data over certain age ranges, or data that is in some way unrepresentative (for example, the accidental inclusion of ill-health retirees can distort the analysis of normal retirees' data at younger ages), or distorted (for example, individual ages in CMI data can occasionally be distorted by multiple records for one individual). Features of this type require careful consideration by the teams responsible for producing the smoothed mortality rates, as different ways of responding to these features can produce significantly different results.

10.7 *Hedge Failure*

10.7.1 There is also the (very important!) issue of hedge failure (or, less extreme, the hedge being less than perfect), the possibility of which can be a major obstacle to the success of any new contract. Hedge failure or basis risk can arise in a number of ways:

- (i) The reference population underlying the security might be different from that of the hedger.
- (ii) The hedger's own mortality experience is stochastic (that is, the hedger can be, and typically will be, exposed to non-systematic longevity risk).
- (iii) Basis risk can arise from trading costs and market illiquidity.
- (iv) Hedge failure can arise from mismatches between the funding requirements of the position being hedged and the hedge itself. For example, if one hedges a forward contract with a futures one, and the underlying then rises in value, the value of the futures position will fall creating margin requirements that have to be met by the end of the day. These funding requirements can be very significant and, if large enough, lead to major actual losses (such as the \$1.3bn loss suffered in such circumstances by Metallgesellschaft in 1994).
- (v) Time basis risk: the risk of choosing the wrong maturity of hedging instrument.

All of these aspects mean that the hedger is unable to hedge his risk perfectly.

10.7.2 The existence of basis risk is important because it affects the

price that a hedger will be prepared to pay for the security. The more basis risk there is, the lower the price. As a consequence, when a new security is under development it is important to assess as far as possible how the intended hedgers might use the security to reduce risk, and to do this it is important to evaluate the degree of basis risk to which they are likely to be exposed. The degree of basis risk needs to be small enough that the hedgers will want to invest in the security to reduce their exposure to longevity risk. If the design of the contract is effective then basis risk will be low for a large number of investors. On the other hand, very low basis risk will often be associated with low liquidity.

11. CREDIT RISK

11.1 *Introduction*

11.1.1 We have described a number of different mortality-linked securities. Each security involves a number of different counterparties and the success of the contract relies on establishing sufficient confidence that contracted cash flows will actually be paid. In some cases such as longevity bonds, these cash flows are as far as 25 years into the future. As an example, in the case of the EIB longevity bond, the mortality swap between BNP Paribas and Partner Re continues for 25 years.

11.1.2 With the EIB longevity bond, investors are directly or indirectly affected by the risk of default by the EIB, BNP Paribas and Partner Re. By far the most significant of these is the risk of default by the EIB, since if either of BNP or Partner Re defaults, the EIB is still contractually obliged to continue payments to the bondholder. However, second order risks exist in the sense that a default by BNP or Partner Re (both of which are AA rated) might, if it happens at the wrong time, trigger a default by the EIB.

11.1.3 In the case of the EIB longevity bond, the investors' main credit risk lies with the EIB itself, but since the EIB is AAA rated it may not be felt necessary for the contract to include a credit-enhancement agreement. However, with other mortality-linked securities, the first point of contact for the investor might be with a lower-rated institution. In such circumstances, it would presumably be essential that a credit-enhancement agreement be put in place: without such an agreement potential investors might be put off subscribing to the issue.

11.1.4 Recall that the primary role of longevity bonds and other mortality-linked securities is to give holders the opportunity to hedge their systematic longevity risks. These investors do not seek exposure to alternative market risks or to additional credit risk. Where any of the key participants carries a risk of default, we therefore also have to consider structures to mitigate this risk. These can make a difference to the price that will be paid and even to the willingness of investors to participate in the first place.

11.1.5 We met one such structure already — the complex ‘background’ details to the EIB bond — in Figure 4. However, there are also many other ways in which the parties concerned might mitigate their credit risk exposures.

11.2 *Credit Insurance*

Credit risk might be mitigated using a credit insurance arrangement, such as a financial guaranty or surety bond: a firm might purchase such insurance to protect itself in the event that its counterparty defaults, or a firm might make itself a more attractive counterparty by purchasing insurance for its counterparty to protect the latter from loss in the event of its own default. However, such insurance can be expensive.

11.3 *Credit Derivatives*

Credit risk might be mitigated using credit derivatives (e.g. credit default swaps) that promise payments if specified credit events occur (e.g. such as the default or downgrading of a counterparty). However, the reference credit events need to be chosen carefully if they are to avoid serious basis risk problems and, as with credit insurance, credit derivatives can be expensive.

11.4 *Credit Enhancement*

Counterparty credit exposures can also be managed using standard credit enhancement methods — collateral agreements, recollateralisation with marking to market (so positions are periodically marked to market, and collateral reassessed accordingly in line with pre-agreed formulas), recouping (in which cash is exchanged when exposures hit pre-agreed limits and payment schedules are re-set to bring the swap value back to zero), credit triggers (in which a counterparty suffering a specified credit downgrade is obliged to close out its swap position and settle its outstanding debts), and mutual termination options (giving either party options to terminate a swap agreement). Each of these methods has proven to be useful in helping firms manage the counterparty credit risks of existing types of swap, and they are especially useful when the swaps are very long-dated ones — as would typically be the case with mortality swaps.

11.5 *Special Purpose Vehicles*

11.5.1 Another commonly used way of mitigating credit risk is through a special purpose vehicle (SPV). One such arrangement (simplified to emphasise the main elements) is illustrated in Figure 7. The initial price paid by the bond holders is supplemented by a further cash injection from the longevity bond issuer with both being paid into an SPV, a company set up especially to service this particular bond. The SPV uses these proceeds to purchase high-grade floating-rate debt, while at the same time enters into an interest-rate swap with an investment bank. The combination of the floating-

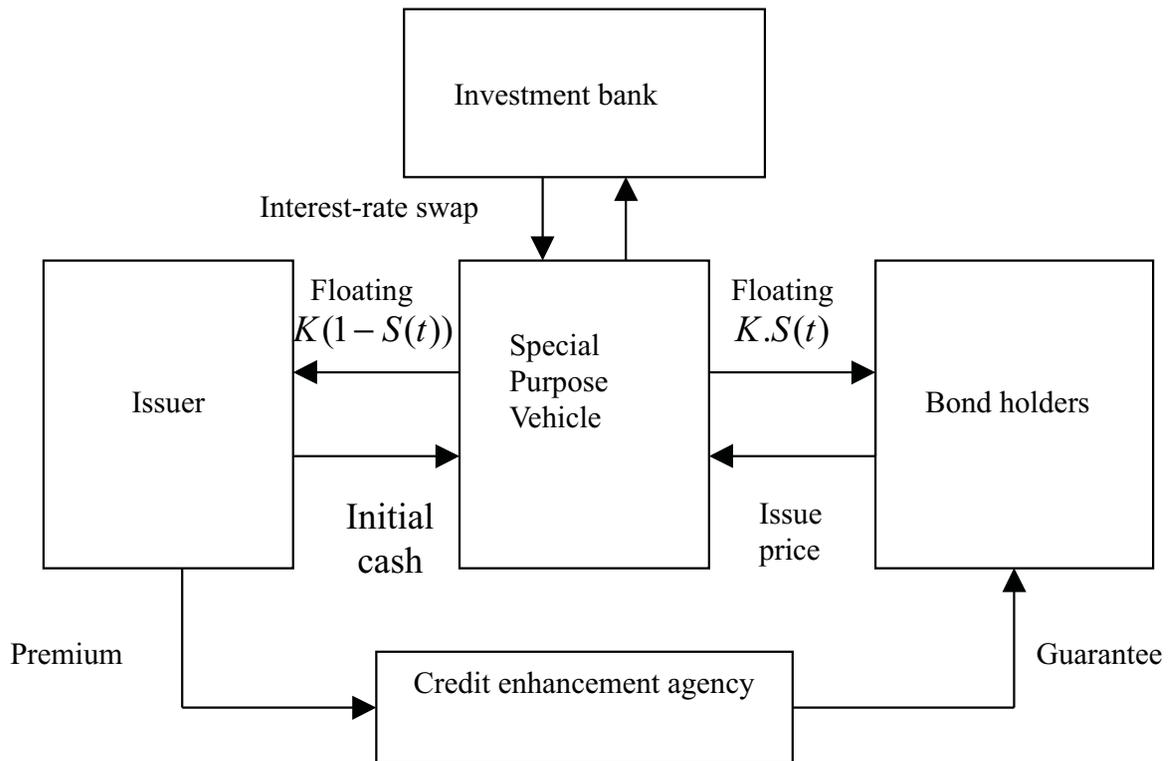


Figure 7. Cash flows in a longevity bond structure involving a special-purpose vehicle

rate debt and the swap will be structured in order to produce cash flows of $\pounds K$ ($\pounds K$ being the initial coupon on the longevity bond) in each future year. These proceeds are then split between the bond holders (who receive the longevity-bond coupons of $\pounds K.S(t)$) and the issuers (who receive the inverse-longevity-bond coupons of $\pounds K(1 - S(t))$). The bond holdings of the SPV and the swap with the investment bank may be subject to a degree of credit risk. If this credit risk is perceived as being too high, then the structure will involve the payment of a premium to a credit enhancement agency which will guarantee to cover any losses to the longevity bond holders should there be any default that would otherwise affect their coupon payments.

11.5.2 The primary role of a credit-enhancement agency is to eliminate so far as possible the exposure of the investor to credit risk by providing insurance that will pay out if any of the counterparties defaults in a way that would otherwise result in a reduction in the future payments to the investor. Credit enhancement agreements might also exist at other points within the overall structure of the security. For example, it might be reasonable in the case of the EIB longevity bond for a credit enhancement agreement to be in place to compensate BNP Paribas or Partner Re should one or the other default on the mortality swap agreement.

11.5.3 The inclusion of a credit-enhancement agency does not of course

completely eliminate the investor's exposure to credit risk, because the agency itself could default. However, if the agency is chosen in an appropriate way then the chances of default by both the credit-enhancement agency and the primary counterparty in the mortality-linked security at the same time should be very small.

12. BARRIERS TO DEVELOPMENT

12.1 In this section, we consider what needs to be done to establish both a spot and a derivatives market in longevity-linked securities.

12.2 The EIB longevity bond did not in fact generate sufficient demand to be launched. It is instructive to look at the possible reasons for this. These can be divided into design issues, pricing issues and institutional issues.

12.3 The following design problems have been suggested, all of which make the bond an imperfect hedge for longevity risk (see, e.g. Barnett-Waddingham, 2004):

- (i) The mortality experience of even large pension funds and life offices may be different from that of the reference U.K. population; the basis risk becomes even worse in the case of small pension funds.
- (ii) The bond only provides a hedge for the longevity risk of males. However, pension funds and life insurers will also be exposed to significant longevity risk from females.
- (iii) Similarly, the bond hedges the longevity risk of 65-year old males. It will therefore be a progressively poorer hedge for younger or older cohorts.
- (iv) The bond only hedges longevity risk for 25 years, handing the tail risk back to the pension fund or life office: around 15% of 65-year old males can expect to live beyond 90.
- (v) The bond gives equal weight to all the lives in the U.K. population, yet the liabilities of pension funds and life offices give greater weight to lives in receipt of larger pensions. Richer pensioners tend to live longer than poorer pensioners.
- (vi) The bond does not hedge inflation risk: it pays out nominal cash flows (adjusted for mortality), but most pensions are uprated in line with RPI, LPI (limited price index) or some other rate.
- (vii) The cash flows are based on the GAD's projection of the probability of death q_x rather than the ONS's central mortality rates $m_x = q_x/(1 - 0.5q_x)$.

12.4 In terms of pricing, the longevity risk premium built into the initial price of the bond was set at 20 basis points. Given that this is the first ever bond brought to market, the markets have no real feeling as to how fair this risk premium is. However, there was some concern that the up-front capital was too large compared with the risks being hedged by the bond (longevity

and interest rate risks), leaving no capital for other risks to be hedged (e.g. inflation).

12.5 The following institutional problems have been raised in the authors' informal discussions with industry practitioners:

- (i) The issue size was too small to create a liquid market.
- (ii) Consultants were reluctant to recommend it to trustees, on the grounds that there is an unwillingness to be the first to leap.
- (iii) Fund managers do not currently have a mandate to manage longevity risk.
- (iv) The number of lives involved in determining cash flows needs to be sufficiently large to avoid anti-gambling provisions which prohibit betting on specific lives.
- (v) There may be a problem with the asymmetric treatment of tax on gains and losses in new hedging instruments: gains might be taxed while losses might not be off-settable against gains elsewhere. Experience shows that persistent lobbying by derivative providers may be necessary to induce the taxing authority to allow hedging transactions to be tax neutral.

12.6 The view has also been expressed (see, e.g. Barnett-Waddingham, 2004) that a key determinant of the future issue of longevity bonds is the availability of sufficient reinsurance capacity. Neither a U.K.-based nor a European Union-based reinsurer was willing to provide cover for the EIB bond, and Partner Re was not prepared to offer cover above the issue size of £540m. This raises the question of whether sufficient reinsurance capacity really exists. One also has to ask whether this capacity problem might be related to the E.U.'s solvency requirements rendering reinsurance cover within the E.U. prohibitively expensive.

12.7 If longevity bonds are to provide effective hedging instruments for the longevity 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. In addition, the problems associated with creating a new liquid market in mortality-linked securities need to be resolved.

12.8 The discussion in Sections 8.2 and 8.3 above suggest important lessons for the development of a mortality-linked futures market:

- (i) The existence of a sufficiently large, active and liquid spot market in longevity bonds with a range of suitable mortality indices is highly desirable before any futures market is started. It is possible for a futures market to start without a traded underlying, but the indices need to be carefully chosen.
- (ii) The mortality index behind the longevity bond or derivative contract must be based on fair, trusted and timely estimates of true mortality and have minimal time basis risk and lag dependence. We might also

note here that the CPI index suffers from the same potential problems, so the survival of the CPI futures contract on the CME suggests that these problems can be overcome.

- (iii) Although mortality indices are calculated infrequently, the spot prices of longevity bonds are likely to exhibit a high degree of volatility on account of the bonds' high duration. However, short-term volatility will be largely due to changes in interest rates rather than changes in mortality, and so there should still be adequate speculator interest.
- (iv) The underlying mortality indices must be few in number, but also appropriately graded. The small number of contracts helps to increase liquidity, but also leads to contemporaneous basis risk, arising from the different mortality experience of the population cohort covered by the mortality index and the cohort relevant to the hedger. The basis risk can be reduced with suitably graded contracts. It is possible in time that basis contracts will develop to hedge basis risk (Wisner, 1997).
- (v) We need to take account of the possibility of natural weaknesses in the market. For example, if we are dealing with 'coupon-based' bonds of which the EIB/BNP bond is an example, there is a potential weak point in the longevity bond market on the supply side, because there are (arguably) few natural issuers. However, the securities themselves can sometimes be structured to get around these problems. For instance, a longevity bond could be structured in, say, a principal-at-risk form (e.g. like the Swiss Re bond) in which case the issuer is the party wishing to hedge the risk. A large U.K. life office told us that depending on circumstances, they could be on either the buy or sell side of the market, so long as the bonds were traded against suitable mortality indices.
- (vi) There is no reason to suppose that liquidity costs in the futures contract would be any higher than for other bond futures contracts.
- (vii) In the early stages of a market, rapid development will be hindered by attempts to protect intellectual property.

13. CONCLUSIONS

13.1 Longevity risk has been around a long time, but its importance has only recently been fully recognised. It has major implications for pensions providers and life insurers, whose ability to make good on their promises depends to a considerable extent on how well they manage this risk. This is of major importance to their policyholders, whose retirements could be ruined if their pension providers fail to deliver. The difficulties of Equitable Life a few years ago demonstrate this danger all too well. The problem of how to live with longevity risk is therefore not going to go away, and is a matter of concern to everyone planning for their retirement.

13.2 Part of the answer is to ensure that annuity providers, insurance companies and pension plans have access to suitable hedge instruments. These are not a substitute for good risk management, of course, but they would of great benefit to well-managed institutions that were serious about the problem of managing their longevity risk exposures.

13.3 The last couple of years have seen efforts to create such instruments. Most particularly, there have been two attempts to offer capital market solutions that enable some of this risk to be laid off. One, the Swiss Re bond (which tackles short-term catastrophic mortality risk), was successfully launched. The other, the EIB bond (which did aim to tackle longevity risk), was withdrawn a year after it was first announced. Despite being highly desirable in principle, it appears that certain features of this bond's design might have discouraged investors. Perhaps the most important of these are excessive basis risk and the fact that the bond requires high upfront capital for the degree of hedging that it provides.

13.4 Fortunately, these design faults are not insurmountable. The problem of basis risk can be overcome by having longer-maturity longevity bonds (such as 'classical' longevity bonds whose payments expire at the same time as the last annuitant in the reference population) or having a range of suitably graded mortality indices (to balance the trade-off between basis risk and liquidity). The problem of upfront costs can be overcome by increasing gearing or by using derivatives, rather than spot market securities. There are various ways such solutions can be implemented, and further work would be needed to identify more fully the trade-offs and choices involved. However, such problems of implementation arise with all new financial instruments, and can — and will — be solved over time.

13.5 Few people doubt that mortality-linked securities are potentially very useful tools for managing longevity risk. Once these teething problems are overcome, the way will be clear for the markets in these securities to develop and mature. We would then be on the cusp of a completely new global financial market in mortality-linked securities. Longevity risk is arguably therefore the next big frontier for financial markets — unless of course someone ruins it all in the meantime and discovers the secret of eternal life.

ACKNOWLEDGMENTS

The authors would like to thank Mark Azzopardi, Tim Cox, David Wilkie, Richard Willets and Andrew Smith, for very helpful comments that have much improved the paper. Data for Figures 5 and 6 were very kindly provided by Aviva through AFPen, the Association of French Pension Funds. The authors are especially grateful to Jean Berthon and Vincent Vandier for their ideas and thoughts on annuity futures. This research was

conducted under Economic and Social Research Council grants RES-000-27-0014 and RES-000-23-1036.

REFERENCES

- BARNETT-WADDINGHAM (2004). Longevity bond to be issued by the EIB. <http://www.barnett-waddingham.co.uk/cms/services/inscomps/news05022/viewDocument>
- BEELDERS, O. & COLAROSSO, D. (2004). Modelling mortality risk with extreme value theory: The case of Swiss Re's mortality-indexed bond. *Global Association of Risk Professionals*, **4** (July/August), 26-30.
- BESSEMBINDER, H. & SEGUIN, P. (1993). Price volatility, trading volume, and market depth: Evidence from futures markets. *Journal of Financial and Quantitative Analysis*, **28**, 21-39.
- BLACK, D.G. (1986). *Success and Failure of Futures Contracts: Theory and Empirical Evidence*. Monograph No. 1986-1, Monograph Series in Finance and Economics, Salomon Brothers Center for the Study of Financial Institutions, Graduate School of Business Administration, New York University.
- BLAKE, D. & BURROWS, W. (2001). Survivor bonds: Helping to hedge mortality risk. *Journal of Risk and Insurance*, **68**, 339-348.
- BLAKE, D., CAIRNS, A.J.G. & DOWD, K. (2003). Pensionmetrics II: Stochastic pension plan design during the distribution phase. *Insurance: Mathematics and Economics*, **33**, 29-47.
- BRORSEN, B. & FOFANA, N. (2001). Success and failure of agricultural futures contracts. *Journal of Agribusiness*, **19**, 129-145.
- CAIRNS, A.J.G., BLAKE, D., DAWSON, P. & DOWD, K. (2005). Pricing the risk on longevity bonds. *Life and Pensions*, **October**, 41-44.
- CAIRNS, A.J.G., BLAKE, D. & DOWD, K. (2004). Pricing frameworks for securitization of mortality risk. In *Proceedings of the 14th International AFIR Colloquium*, Boston, November 2004, pp509-540.
- CAIRNS, A.J.G., BLAKE, D. & DOWD, K. (2005a). A two-factor model for stochastic mortality with parameter uncertainty. To appear in *Journal of Risk and Insurance*.
- CAIRNS, A.J.G., BLAKE, D. & DOWD, K. (2005b). Pricing death: Frameworks for the valuation and securitisation of mortality risk. Working paper, Heriot-Watt University.
- CARLTON, D.W. (1984). Futures markets: Their purpose, their history, their growth, their successes and failures. *Journal of Futures Markets*, **4**, 237-271.
- CMI (2005). Projecting future mortality: Towards a proposal for a stochastic methodology. Working Paper 15 of the Continuous Mortality Investigation. The Faculty of Actuaries and Institute of Actuaries.
- CORKISH, J., HOLLAND, A. & VILA, A.F. (1997). The determinants of successful financial innovation: An empirical analysis of futures innovation on LIFFE. London: Bank of England.
- COWLEY, A. & CUMMINS, J.D. (2005). Securitization of life insurance assets and liabilities. *Journal of Risk and Insurance*, **72**, 193-226.
- COX, S.H. & LIN, Y. (2004). Natural hedging of life and annuity mortality risks. In *Proceedings of the 14th International AFIR Colloquium*, Boston, pp483-507.
- CURRIE, I.D., DURBAN, M. & EILERS, P.H.C. (2004). Smoothing and forecasting mortality rates. *Statistical Modelling*, **4**, 279-298.
- DOWD, K., BLAKE, D., CAIRNS, A.J.G. & DAWSON, P.E. (2005). Survivor swaps. To appear in *Journal of Risk and Insurance*.
- DOWD, K., CAIRNS, A.J.G. & BLAKE, D. (2005). Mortality-dependent measures of financial risk. To appear in *Insurance: Mathematics and Economics*.
- EDERINGTON, L.H. (1979). The hedging performance of the new futures markets. *Journal of Finance*, **39**, 157-170.

- FORFAR, D.O., McCUTCHEON, J.J. & WILKIE, A.D. (1988). On graduation by mathematical formula. *Journal of the Institute of Actuaries*, **115**, 1-149.
- GRAY, R.W. (1978). Why does futures trading succeed or fail? An analysis of selected commodities. In A.E. Peck (ed.) *Readings in Futures Markets: Views from the Trade*. Chicago: Chicago Board of Trade.
- LIN, Y. & COX, S.H. (2005). Securitization of mortality risks in life annuities. *Journal of Risk and Insurance*, **72**, 227-252.
- OLIVIER, P. & JEFFERY, T. (2004). Stochastic mortality models. Presentation to the Society of Actuaries of Ireland: see http://www.actuaries.ie/Resources/events_papers/PastCalendarListing.htm
- PATEL, K. (1994). Lessons from the FOX residential property futures and mortgage interest rate futures market. *Housing Policy Debate*, **5**, 343-360.
- PIEROG, K. & STEIN, J. (1989). New contracts: What makes them fly or fail? *Futures*, **September**, 51-54.
- ROHRER, M.B. (2004). The relevance of basis risk in the weather derivatives market. In M. Constantino and C.A. Brebbia (eds.). *Computational Finance and its Applications*. Southampton: WIT Press.
- SRINIVASAN, S. (2004). CPI futures at CME. Chicago: Chicago Mercantile Exchange.
- WADSWORTH, M., FINDLATER, A. & BOARDMAN, T. (2001). Reinventing annuities. Paper presented to the Staple Inn Actuarial Society. <http://www.sias.org.uk/>
- WISNER, R. (1997). Understanding risk in basis contracts. Working paper, Department of Economics, Iowa State University.