



City Research Online

City, University of London Institutional Repository

Citation: Blake, D., Boardman, T. and Cairns, A. J. G. (2014). Sharing Longevity Risk: Why Governments Should Issue Longevity Bonds. *North American Actuarial Journal*, 18(1), pp. 258-277. doi: 10.1080/10920277.2014.883229

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/6842/>

Link to published version: <http://dx.doi.org/10.1080/10920277.2014.883229>

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.



DISCUSSION PAPER PI-1002

Sharing Longevity Risk: Why Governments Should Issue Longevity Bonds

David Blake, Tom Boardman and Andrew Cairns

February 2013

ISSN 1367-580X

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

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

**Sharing Longevity Risk:
Why Governments Should Issue Longevity Bonds**

David Blake,
Pensions Institute
Cass Business School
106 Bunhill Row
London, EC1R 1XW
United Kingdom
(d.blake@city.ac.uk)

Tom Boardman,
Visiting Professor, Pensions Institute
Cass Business School
106 Bunhill Row
London, EC1R 1XW
United Kingdom
(tvboardman@gmail.com)

Andrew Cairns,
Maxwell Institute for Mathematical Sciences and Department of Actuarial
Mathematics and Statistics
Heriot-Watt University
Edinburgh, EH14 4AS
United Kingdom
(A.J.G.Cairns@ma.hw.ac.uk)

18 February 2013

Abstract

Government-issued longevity bonds would allow longevity risk to be shared efficiently and fairly between generations. In exchange for paying a longevity risk premium, the current generation of retirees can look to future generations to hedge their systematic longevity risk. Longevity bonds will lead to a more secure pension savings market, together with a more efficient annuity market. By issuing longevity bonds, governments can aid the establishment of reliable longevity indices and key price points on the longevity risk term structure and help the emerging capital market in longevity-linked instruments to build on this term structure with liquid longevity derivatives.

Key words: Longevity risk, longevity bonds, public policy, political economy

JEL classifications: G22, G23, G24, G28, H11, H63, J11, J18

I. Introduction

Longevity bonds pay declining coupons linked to the survivorship of a cohort of the population, say 65-year-old males; for example, the coupon payable at age 75 (i.e., 10 years after the issue date of the bond) will depend on the proportion of 65-year-old males who survive to age 75; they have no principal repayment. They are designed to hedge systematic (also known as aggregate or trend) longevity risk.

Insurance companies and pension plan providers face the risk that retirees might on average live longer than expected. Longevity risk is a substantial risk that might adversely affect both the willingness and ability of financial institutions to supply retired households with financial products to manage wealth decumulation in retirement. In this paper, we explain how governments issuing longevity bonds can act as a catalyst to facilitate the transfer of a proportion of this risk to the capital markets. We highlight the benefits that would flow from a transparent and liquid capital market in longevity risk, and we argue that there is an important role for governments to play in helping this emerging market to grow. We also show how the government might consider how to price longevity bonds in the face of potential demand from defined benefit (DB) and defined contribution (DC) plans and from annuity providers. Our line of reasoning comes from working in the UK, but we believe that what we argue here has validity for all countries with mature funded pension systems.

The UK pension fund industry is the second largest in the world by value, with assets of around 20% of those held in the USA. However, the UK lifetime annuity market is much larger than in the US – around 500,000 annuities are set up each year at a cost of £12bn, mainly as a result of the effective requirement to buy life annuities as part of DC pension plan provision.


A well-functioning annuity market will become increasingly important as DC plans mature, not just in the UK, but in all countries where DC pension provision becomes the norm. The importance of DC pensions and, in turn, lifetime annuities is growing rapidly as governments cut social security pensions and companies move away from DB plans. DC plans *have to* work effectively if people are going to be prepared to

save privately for their pensions. However, a growing weakness in DC plans is the inability of annuity providers to hedge the systematic longevity risk they face. Systematic longevity risk might affect the price and availability of annuities, as well as insurance company solvency. Every country with DC pension plans will sooner or later have to confront the problem of dealing with systematic longevity risk.

We therefore believe that the time is right for governments to set up a working party to undertake a cost-benefit analysis of the government issuance of longevity bonds.

II. What is longevity risk?

Figure 1: Decomposition of longevity risk



1

Figure 1 shows that longevity risk is driven by two underlying risks: random variation risk and trend risk. Random variation risk is the risk that individual mortality rates differ from the outcome expected as a result of chance – some people will die before their life expectancy, some will die after.¹ Trend risk is the risk that unanticipated changes in life-style behaviour or medical advances significantly improve longevity.²

¹ The mortality rate for a given age measures the frequency of occurrence of deaths of people of the given age in a defined population during a specified time interval, typically one year. Mortality rates are derived from crude death rates which are calculated as the ratio of deaths to the exposed population, i.e., the number of lives at the start of the period exposed to the risk of dying during a specified time interval, typically one year. A survivor (or survival) rate for a given age measures the proportion of people of the given age surviving a specified time interval. The survivor rate at age 65 equals (1 –

Private-sector institutions can deal with a ‘specific risk’ like random variation risk by pooling and relying on the law of large numbers to reduce the variability of this risk. Trend risk, on the other hand, is, like inflation risk, a ‘systematic risk’ that cannot be diversified away by pooling³ and, indeed, the more business an insurer pools, the bigger the relative impact of trend risk. The private sector is unable to hedge this risk effectively without a suitable hedging instrument. We will argue that there is a key role for governments to help the private sector by issuing longevity bonds – particularly by issuing bonds that provide ‘tail risk’ protection against trend risk – and by helping with the construction of national longevity indices.

III. Why should we be concerned about longevity risk and who bears it?

Longevity risk is borne by every institution making payments that depend on how long individuals are going to live. These include DB pension plan sponsors, insurance companies selling life annuities and governments through the social security pension system and the final salary pension plans of public-sector employees. The situation is particularly acute for insurance companies operating in the European Union (EU) where a new regulatory regime, Solvency II, is due to be introduced in 2014.⁴ The current Solvency II proposals, if adopted, will require insurers to hold significant additional capital to back their annuity liabilities if longevity risk cannot be hedged effectively or marked to market.

mortality rate at age 65). Life expectancy measures the average number of years a person of a given age would live under a given set of mortality conditions. Life expectancy is usually computed on the basis of a life table showing the probability of dying at each age for a given population according to the age-specific death rates prevailing during a specified period. For example, life expectancy at 65 = $0.5 + (1-q(65)) + (1-q(65))*(1-q(66)) + (1-q(65))*(1-q(66))*(1-q(67)) + \dots + (1-q(65))* \dots *(1-q(120))$ and $q(120)$ is typically set to unity and $q(65)$ is the mortality rate at age 65, etc. We also need to distinguish between period life expectancy which makes no allowance for future improvements in mortality rates – and so assumes, for example, that $q(67)$ in the above formula will equal the mortality rate of today’s 67-year-olds – and cohort life expectancy which makes such an allowance – and hence will involve a lower $q(67)$ than used to calculate period life expectancy.

² Factors such as obesity and environmental degradation could eventually lead to a trend decline in life expectancy.

³ Milevsky et al. (2006) prove this result.

⁴ See Appendix A for more details about Solvency II.

By any measure, longevity risk is a significant risk. Global private-sector pension liabilities are of the order of \$25trn.⁵ In the UK alone, private-sector DB pension liabilities equal £1,340bn, while DC pension assets amount to £737bn (including £150bn in annuities with insurance companies).⁶ It has been estimated that every additional year of life expectancy at age 65 adds around 3 percent or £33bn to the present value of DB pension liabilities in the UK, with a similar impact on lifetime annuities.⁷ The most recent estimates for UK state pension liabilities were £3,843bn in respect of social security pensions, £852bn in respect of the unfunded pension plans of public-sector employees, and £313bn in respect of the funded plans of public-sector employees (principally local government employees).⁸ This implies that UK government-backed longevity-linked liabilities exceed £5trn.⁹

In addition to being extensive, longevity risk in the private sector is beginning to become concentrated, especially in the UK. Private-sector companies in the UK are moving rapidly away from DB pension provision. They are beginning to offload the legacy longevity risk that they still hold either by buying-in annuities from life companies to cover their pensions-in-payment or by undertaking bulk buy-outs of their liabilities, again with life companies.^{10,11} In providing these indemnification solutions for DB pension plans, insurance companies are beginning to play a big role in aggregating longevity risk in the economy.

The DB plans in private-sector companies in the UK are being replaced with occupational DC plans – the equivalent of 401(k) plans in the USA – and, in so doing,

⁵ OECD (2011) and Life and Longevity Market Association

⁶ Levy (2012) and Association of British Insurers; the figures are for end-2010.

⁷ Pension Protection Fund and the Pensions Regulator (2006, Table 5.6).

⁸ Hobbs (2012); the figures are for end-2010.

⁹ The UK government has linked the social security pension age to increases in life expectancy and is planning to do the same for public sector employees, so this figure is not expected to increase in future as it has in the past.

¹⁰ Bulk-buyouts transfer the pension liabilities in corporate pension plans to insurance companies. This market began in earnest in the UK in 1999, when the Prudential Assurance Company did £1bn of business.

¹¹ There is also an increasing use of longevity swaps provided by both insurance companies and investment banks (Hymans Robertson, *Buy-outs, Buy-ins and Longevity Hedging* (various issues)). A longevity swap exchanges fixed for floating survivor rates over the tenor of the swap. The fixed rates might be set equal to the expected rates in Figure 2 below plus the longevity risk premium. The floating rates are the realized rates which could be above or below the fixed rate. Each year, the pension plan or annuity provider pays the fixed rate and receives the floating rate and thereby locks in the cost of the pension or annuity payments. The first suggestion for longevity swaps – or survivor swaps – was made in Dowd et al. (2006).

companies are passing the longevity risk back to their employees. So individuals should be concerned because there is a real risk that they will outlive their wealth – this is the specific risk identified in Figure 1 – if they do not hedge this risk by buying life annuities. In countries such as the UK and Chile where annuitization of DC pension pots is either mandatory or strongly incentivized, it will again be life companies that provide these annuities.

So all the trends in pension provision – increasing demand from DB plans to use annuities to back their pensions in payment, the growing demand from DB plans for bulk buy-outs, the overall growth in both the number and size of DC pension funds and the associated growth in the number of pensioners with DC funds reaching retirement – are pointing to a big increase in demand for annuities provided by insurance companies.

There are two problems associated with this increased demand. First, there is the danger that this could result in an unhealthy concentration of risk amongst a small number of insurance companies. Second, there is insufficient capital in the insurance/reinsurance industry to deal with total global private-sector longevity risk.

Under Solvency II, it is proposed that insurance liabilities are increased by the addition of a market value margin (MVM) reflecting the cost of capital to cover ‘non-hedgeable’ risks. For annuity companies this is principally longevity risk. It is currently proposed that in the absence of a hedging instrument for longevity risk, EU insurers will have to charge a 6% cost of capital above the risk-free rate when calculating the MVM. As a consequence of the long-dated nature of annuities, this calculation could result in the amount of capital held for longevity risk approximately doubling from current levels. The resultant extra capital for longevity risk and other Solvency II impacts¹² would have to be passed on to customers and the money’s worth of annuities could fall by up to 10%.¹³

¹² For example, the loss of upfront allowances for the liquidity premium and for credit risk.

¹³ Tully (2011). Of this 10%, industry insiders estimate that 7% is accounted for by the lost allowances for the liquidity premium and for credit risk, with the remaining 3% due to the absence of a longevity risk hedge. With £12bn annual sales of annuities in the UK, this implies a cost to every new annual cohort of retirees in the UK alone of £360mn.

The only realistic way of handling the issues of concentration and sufficient capital, is to find an efficient way or passing some of the risk onto governments and the capital markets. The alternative is poorer value annuities, an annuity market prone to insolvency or, in the extreme, no private-sector annuity market at all. All governments that have encouraged the growth of DC pension provision should be concerned about this. But, by issuing longevity bonds, governments can help to overcome these problems.

IV. How can longevity bonds hedge systematic longevity risk?


In order to see how a longevity bond can hedge systematic longevity risk, we need to both quantify longevity risk and identify where it is concentrated. Figure 2 presents a survivor fan chart¹⁴ derived using the Cairns-Blake-Dowd (CBD) stochastic mortality model.¹⁵ The fan chart shows the uncertainty surrounding projections of the number of survivors to each age from the cohort of males from the national population of England and Wales who are aged 65 at the end of 2006.¹⁶ The bars indicate the 90% confidence interval on the projected survivor rate for each age out to 115. The line in the middle of each bar indicates the expected proportion of the cohort to survive to each age. The Figure shows that there is little uncertainty out to age 75: we can be fairly confident that approximately 19% will have died by 75. The uncertainty peaks at age 93: the confidence interval band is widest at this age. The best estimate is that 36% will survive to age 90, but it could be anywhere between 30% and 41%. This is a very large range. The Figure also shows the extent of the so-called ‘tail risk’ after age 90: there is some probability – even if small – that some members of this cohort will live beyond 110.

¹⁴ Blake et al. (2008).

¹⁵ Cairns et al. (2006). This model is briefly explained in Appendix B.

¹⁶ The CBD model was estimated using data between 1991 and 2006. The historical period over which a stochastic mortality model such as the CBD model is estimated is certainly important for both getting a good fix on the future trend improvements in mortality rates and on their volatility around this trend. However, this does not necessarily mean that a longer data period is better. If there has been a significant change in the trend, then this suggests the model should be estimated over a short period for the purpose of getting a reliable estimate of the latest trend. On the other hand, a longer period might be used to get an estimate of long-run volatility. This is a matter of experimentation. The results we present here are purely illustrative, although they were compared for with consistency with the official Office for National Statistics 2008 projections. Much more analytical work would have to be done using a wider range of models before a real-world longevity bond could be issued.

Figure 2: Survivor fan chart - Males aged 65



Note: Derived from the Cairns-Blake-Dowd stochastic mortality model, estimated on English and Welsh male mortality data for 65-year olds over the period 1991-2006

A survivor fan chart is very useful to a pension plan or annuity provider since it shows the likely range of pensioners or annuitants from a given birth cohort surviving to each age. If more survive to each age than was expected, the pension plan or annuity provider has to make higher total pension or annuity payments than was anticipated. The opposite holds if fewer survive to each age than was anticipated. The best estimate expectation of life is 20.5 years; the 5% confidence level expectation is 19.4 years and the 95% confidence level expectation is 21.8 years.

We will now show how a longevity bond with the following characteristics can help to hedge systematic longevity risk:

- The bond pays coupons that decline over time in line with the actual mortality experience of a cohort of the population, say 65-year-old males from the national population: so the coupons payable at age 75, for example, will depend on the proportion of 65-year-old males who survive to age 75.
- Coupon payments are not made for ages for which longevity risk is low: so, for example, the first coupon might not be paid until the cohort reaches age 75 (such a bond would be denoted as a deferred longevity bond).
- The coupon payments continue until the maturity date of the bond which might, for example, be 40 years after the issue date when the cohort of males reaches age 105.

- The final coupon incorporates a terminal payment equal to the discounted value of the sum of the post-105 survivor rates to account for those who survive beyond age 105. The terminal payment is calculated on the maturity date of the bond and will depend on the numbers of the cohort still alive at that time and projections of their remaining survivorship. It is intended to avoid the payment of trivial sums at very high ages.
- The bond pays coupons only and has no principal repayment.

Figure 3: Deferred Longevity Bond for male aged 65 with 10-year deferment




Figure 3 shows the possible range of coupon payments on a deferred longevity bond based on the national population of English and Welsh males who were aged 65 at the end of 2006. Such a bond would provide a hedge for the systematic longevity risk faced by pension plans and annuity providers. If population survivorship is higher at each age than was expected, the bond pays out higher coupons. This is what pension plans and annuity providers need to help match the higher than expected pensions and annuity payments they need to make. If, on the other hand, survivorship is lower at each age than was expected, the bond pays out lower coupons. But the pension plans and annuity providers are not likely to mind this, since their pensions and annuity payments are also likely to be lower.

However, it is important to recognize that the bond will only provide a perfect hedge for the systematic longevity risk faced by pension plans and annuity providers if the plan members and annuitants have exactly the same mortality experience over time as the cohort underlying the bond. If the plan members and annuitants have a mortality experience that differs from that of the national population, this will introduce basis risk.¹⁷ In practice, there will always be some basis risk. One reason for this is that pension plans and annuity books have far fewer members than the national population and will therefore experience greater random variation risk than the national population and this is likely to cause the mortality experience of a sub-population to diverge from that of the national population over time, even if they have the same mortality profile at the outset.

Another reason is that most pension plans and annuity books will not have the same mortality profile as the national population, even to begin with. There can be differences in age, gender and socio-economic composition. Different birth cohorts have different survivor rates to each age. While survivor rates to each age tend to increase over time, in line with the trend improvement in longevity, they do not do so uniformly: some birth cohorts experience faster improvements than others.¹⁸ Females, on average, live longer than males. Professionals tend to live longer than white-collar workers who in turn tend to live longer than blue-collar and manual workers. But it is not simply the differences in life expectancies between these various groups that are important, it is unexpected changes in the trends in their survivorship experience that causes basis risk.

Yet another reason for basis risk involves the difference between ‘lives’ and ‘amounts’. A population longevity index¹⁹ will weight each life equally, but members of the higher socio-economic groups will tend to have higher pensions and annuities than members of the lower socio-economic groups. They are also more likely to have multiple annuities. The directors of a small manufacturing company are likely to

¹⁷ This is the risk that the ‘underlying’ – in this case, the survivor rates of the particular population being hedged – does not move in line with the hedging instrument – which, in this case, depends on the survivor rates of the national population.

¹⁸ Willetts (2004), Richards et al. (2006).

¹⁹ This is an index based on the mortality experience of the national population.

represent a large share of the company's pension plan liabilities and are more likely to live longer than the average member. All these factors will increase basis risk and its complexity.

In theory, there could be a longevity bond for both males and females, for each age and for each socio-economic group. Such granularity of the longevity bond market would allow a high degree of hedge effectiveness to be achieved. But it would also result in negligible liquidity or pricing transparency: the more bonds there are, the less trading there will be in each bond and the less frequently the bonds will be priced. As is the case in other markets – especially derivatives markets – a small number of suitably designed bonds should provide an appropriate balance between hedge effectiveness, liquidity and pricing transparency.²⁰

Not only are longevity bonds useful for hedging systematic longevity risk once pensioners have retired, they could be used to hedge systematic longevity risk and long-term investment risk in the period leading up to retirement. A typical DC plan will use a life-style (or life-cycle) investment strategy. This involves a high weighting in equities and other growth assets in the early stages of the accumulation process in order to benefit from the equity risk premium. There is then a systematic switch to less volatile assets, typically long-dated fixed-income bonds, during the final stages of the accumulation process – the so-called glide path to retirement – in order to reduce the volatility of the lifetime retirement income secured at retirement. While the fixed-income bonds hedge the interest-rate risk in the purchase of an annuity,²¹ they do not hedge the longevity risk.²²

Both interest-rate risk and longevity risk could be hedged along the glide path if plan members invested in a fund containing longevity bonds. This would give plan members greater certainty of income in the run up to retirement. This follows because the price of future lifetime annuities (at the member's retirement date) should be

²⁰ See the discussion in section 8 of Blake et al. (2006).

²¹ Since annuity providers buy bonds to make the annuity payments, annuities are subject to interest-rate risk. If interest rates fall, bond prices rise and this will reduce the amount of the annuity that can be paid from a given lump sum.

²² If longevity improves at a higher rate than that expected along the glide path, this too will reduce the amount of the annuity that can be paid from a given lump sum.

highly correlated with the value of this fund which will rise if longevity improves faster than expected or if long-term interest rates fall, and reduce if longevity expectations decline or interest rates rise. The fund might be a better way of providing income security from a DC pension plan at retirement than the alternative of purchasing deferred annuities, since the annuity provider might have to hold significant capital against the deferred annuities it sold (at least this is true in the UK), the cost of which would have to be passed onto the member.

V. Why should the government issue longevity bonds?

In principle, longevity bonds could be issued by private-sector organizations. It has been argued that pharmaceutical companies would be natural issuers, since their revenues are positively linked to survivorship: the longer people live, the more they will spend on medicines.²³ While this is true, the scale of the demand for longevity bonds far exceeds conceivable private-sector supply from companies such as pharmaceuticals. Further, there would be significant credit risk associated with the private-sector issuance of an instrument intended to hedge a systematic risk many years into the future. In practice, we believe that the only realistic issuer of longevity bonds in scale is the government.^{24,25}

We believe that there are three important reasons why the government should engage in sharing longevity risk with the private sector. It:

- has an interest in ensuring there is an efficient annuity market
- has an interest in ensuring there is an efficient capital market for longevity risk transfers
- is best placed to engage in intergenerational risk sharing, such as by providing tail risk protection against systematic trend risk.²⁶

²³ Dowd (2003).

²⁴ The first suggestion for governments to do this was made in Blake and Burrows (2001).

²⁵ See section X below for a critique of this view.

²⁶ See Bohn (2012) for a formal model of intergenerational risk sharing in the face of shocks to labour productivity, return on capital and longevity. Bohn recommends governments should issue both wage-

A. An efficient annuity market for pensioners

The government has an interest in ensuring there is an efficient annuity market, given its desire to encourage retirement savings in DC pension plans that rely on annuities to turn pension savings into guaranteed lifetime retirement income. If the private sector is unable to hedge systematic longevity risk, it increases the likelihood that insurance companies stop selling annuities or increase annuity prices which would reduce pensioner income in retirement.

A consequence of the above is that governments might find themselves having to pay additional means-tested benefits to supplement pensioners' incomes, as well as receiving lower income tax and expenditure taxes (such as value added tax in the UK) from pensioners due to their lower incomes.²⁷ This will, *ceteris paribus*, lead to higher taxes on the working population. This outcome will therefore not be popular with workers or pensioners. Further, workers are likely to reduce savings into DC pension plans. Those that do continue to save in DC plans will face even greater uncertainty about their prospective pension income, since an efficient private-sector annuity market might no longer be in existence when they retire.

B. An efficient capital market for longevity risk transfers

The capital markets have a key role to help ensure there is an efficient annuity market and to reduce concentration risk. It can therefore also be argued that the government has an interest in ensuring there is an efficient capital market for longevity risk transfers. There are two areas where government support is required.

First, the government can help with the construction of national longevity indices. It is for reasons of accuracy that longevity indices would most likely have to be based on national mortality data. A key component of the success of the new capital market

and longevity-indexed bonds, since these would help to reduce both the mismatch between pension assets and liabilities and the pension fund's dependence on corporate sponsors.

²⁷ Many of the people buying annuities in the UK are also on means-tested benefits. Any reduction in annuity payments arising from more onerous capital requirements resulting from insurers being unable to hedge longevity risk will *immediately* increase means-tested benefits.

will be the timely publication of accurate and independently calculated longevity indices. The longevity indices would cover mortality rates, survivor rates and life expectancies for both males and females.

Only the government has access to the information necessary to produce these indices on account of the legal requirement to report deaths and related information such as dates of death and birth and gender to an official agency, which in the UK is the General Register Office of Births, Marriages and Deaths.²⁸ Further, only the government has access to the information needed to estimate the size of the exposed population. In the UK, this is currently derived from decadal censuses with annual updates between censuses based on reported deaths and estimated migration flows. However, the resulting estimates are not accurate enough at high ages. It is important to be able to track a cohort over time, particularly at high ages: the government is in a unique position to do this, since it makes social security pension payments to almost every old person and needs to keep good records to do this. While longevity indices based on social class would be useful, the social class of a deceased person is not recorded at the time of death and while attempts have been made to construct social class indices, based on factors such as zip code or post code, these lack the accuracy of national indices. A similar argument would hold for longevity indices based on amounts rather than lives.²⁹

Second, the government can make an important contribution by issuing longevity bonds to facilitate price discovery, thereby encouraging capital market development. Longevity risk is not currently actively traded in the capital markets, so we do not have a good estimate of its market price or premium.³⁰ But if the government issued a small number of longevity bonds, this would help to establish and maintain the market-clearing 'price points' for longevity risk at key ages and future dates, and hence establish a market price for longevity risk. In other words, the bonds would

²⁸ The government will always have more refined information than the private sector as a result of data protection legislation. This legislation prevents the release of information that would allow an individual – even one who has died – to be identified. Mortality data will only be published in a sufficiently aggregated form – in terms of date and location of death – that makes it impossible for specific individuals to be identified.

²⁹ For an examination of longevity hedging using longevity indices, see Coughlan et al. (2011).

³⁰ The longevity risk premium is paid by the longevity bond's buyer to the bond's issuer to remove systematic longevity risk. It therefore results in a lower coupon that the bond's issuer has to pay the bond's buyer for purchasing the bond, thereby lowering the effective yield on the bond.


help to establish the riskless term structure for survivor rates for ages above 65 for future years.³¹ There is a clear analogy with the fixed-income and index-linked (TIPS in the US) bond markets. In these markets, the issue of government bonds helped to establish the riskless term structures for interest rates and inflation rate expectations, respectively, for terms out to 50 years or more. The private sector was then able to issue corporate fixed-income and index-linked bonds with different credit risks (AAA, AA, etc.) and establish credit term structures above the riskless benchmark curves.

The longevity risk term structure is more complex than either the interest rate or inflation term structures, since it is two-dimensional – involving age as well as time – whereas the latter are one-dimensional, involving only time. The longevity risk term structure is therefore a two-dimensional surface, rather than a line: cohorts move diagonally across the surface over time, getting one year older with every passing year, with some members of the cohort dying each year. This is demonstrated in Figure 4 which shows the cash flows on two deferred longevity bonds: one bond based on male lives from the national population aged 65 and one bond based on male lives from the national population aged 75. Each bond is specified by four dates: the birth year of the cohort being tracked (e.g., 1945), the issue date (e.g. 2010), the first payment date (e.g., 2020) and the last payment date (e.g., 2050).³² There is a corresponding mortality term structure for females, so longevity bonds are also identified by gender (M or F).

³¹ Currently, the survivor rates for future years are based on model projections, such as the CDB model. Figure 2 illustrates this for males aged 65 at the end of 2006. The theoretically fair price of a longevity bond could therefore be determined using the CBD model. However, with a traded market in longevity bonds, a market view of future survival rates would replace model projections and the resulting price points would be used in determining the market price of the bonds. Pricing-to-market would replace pricing-to-model.

³² If a strips market in longevity bonds develops – as happens with fixed-income and index-linked bonds – then hedgers could buy the subset of the coupon payments that most closely meets their hedging requirements, rather than having to buy the whole bond. In addition, if the individual coupons in Figure 4 are traded separately, this will allow more accurate determination of the price points for longevity risk along the diagonals of the longevity risk term structure.

Figure 4: Longevity Bond cash flows across ages and time



The establishment of a market price for longevity risk would be particularly useful for EU insurance companies operating under Solvency II. The maximum longevity risk premium that an annuity provider would be willing to pay to buy a longevity bond would be related to the level of capital that the regulators agree can be released as a result of holding the longevity bond to back annuity liabilities.³³

The establishment of price points will also help to facilitate the capital market development of longevity swaps and other longevity derivatives similar to the interest-rate and inflation swaps that developed in the fixed-income and index-linked bond markets. Market participants were able to use market interest-rate and inflation expectations rather than projections from models. The same would happen in the longevity swaps market. The longevity swaps market began to develop in the UK in 2007-09 with eight publicly announced swaps involving six annuity providers and two pension funds. A number of global investment banks and reinsurers intermediated the deals – J.P. Morgan, Deutsche Bank, RBS, Credit Suisse, Goldman Sachs and SwissRe – and the longevity risk was passed through to investors – such as insurance-

³³ It will also be related to the extent of the basis risk that remains unhedged and potentially the size of any illiquidity premium contained in the price of longevity bonds. If longevity bonds are not actively traded, investors will demand an illiquidity premium to hold them and the regulator might be reluctant to accept that the bonds' prices can be used for mark-to-market pricing for capital release purposes.

linked securities (ILS) investors, hedge funds, sovereign wealth funds, family offices and endowments – attracted by a new asset class that is uncorrelated with traditional asset classes, such as equities, bonds and real estate.


C. Intergenerational risk sharing

The government is the only agency in society that can engage in intergenerational risk sharing on a large scale and enforce intergenerational contracts.³⁴ This is important, given that longevity risk is a risk that crosses a number of generations.

This is how the intergenerational risk sharing operates. The government would receive a longevity risk premium by issuing longevity bonds. In effect, the current retired population pays future generations an insurance premium to hedge its systematic longevity risk. If, in equilibrium, the risk premium is sufficient to ensure that the generation bearing the risk is adequately compensated, then each generation is treated fairly. The current generation of pensioners derives benefit from annuity companies being able to use government-issued longevity bonds to provide better value annuities. The premium that this generation pays for taking away the longevity risk is effectively the premium required to compensate the younger generations to whom the government is passing on the risk in the form of possible higher taxes to enable the government to continue paying pensions to members of the current generation who live longer than expected.

³⁴ In the private sector, long-term contracts can involve significant credit risk as mentioned above and collateralization can introduce significant frictional costs

Figure 5: Deferred Tail Longevity Bond for male aged 65



A key role for government in this context is to provide a hedge for systematic longevity risk by offering tail risk protection against trend risk. Once the market for longevity bonds has matured, in the sense of producing stable and reliable price points in the age range 65-90, the capital markets can take over responsibility for providing the necessary hedging capacity in this age range using longevity securities and derivatives. All that might then be needed would be for the government to provide a continuous supply of deferred tail longevity bonds with payments starting from age 90 in order to allow pension plans and insurers to hedge their tail risk.³⁵ Figure 5 illustrates the cash flows on such a bond. These bonds will be necessary on a permanent basis, since the capital that annuity providers would be required by the regulator to post in order to cover this risk would be very high in the absence of a close matching asset. The bonds are also necessary because the investors who have recently become interested in taking the other side of the longevity swaps market have no appetite for hedging long-duration tail longevity risk.

³⁵ Pension plans and annuity providers might still be willing to invest in government-issued longevity bonds covering the age range 65-90 if they are competitively priced compared with capital market hedges.

VI. What is the potential demand for longevity bonds?

The demand for longevity bonds is driven principally by the growth of DC pensions and the growing maturity of DB plans. The market in DB longevity risk management is new and there is a significant programme currently being implemented in the UK by investment banks and actuarial consultants to educate DB pension plan trustees and annuity providers about the benefits of longevity risk hedging. Although the investment banks have an incentive to talk up the market, the demand is genuine. We believe that the potential demand for longevity bonds is substantial.

In the UK alone: of the £1.3trn in DB private-sector pension liabilities, around £600bn relate to pensions in payment; of the approximately £600bn in accumulated DC pension assets, £200bn relate to people over age 55; and insurance companies are committed to making annuity payments valued in excess of £150bn.

We believe that a suitable initial issuance of longevity bonds (with 10-year deferment) by the UK government could be four bonds: LBM(65,75), LBF(65,75), LBM(75,85) and LBF(75,85).³⁶ The size of each bond issue will depend, in part, on price and this will be considered in the next section. However, the total issuance is likely to be small in relation to the overall size of the government bond market and is unlikely to become a principal funding source for government.³⁷ Nevertheless, the issuance will have significant value, since it will improve the efficiency of the annuity market as well as providing a useful risk management tool for DB plans.

VII. Pricing considerations

Ultimately, the demand for longevity bonds will depend on their price. Demand will be higher the closer the government offers the bonds at true economic cost, i.e., charges a fair, but not excessive, longevity risk premium. It is right that the government seeks to charge a fair risk premium on longevity bonds because this

³⁶ LBM(65,75) is a longevity bond for males aged 65, with the first coupon paid at age 75, etc.

³⁷ Total UK government bond issuance will exceed £700bn over 5 years as a consequence of the fallout from the 2007-08 Global Financial Crisis.

ensures intergenerational fairness. The expected cost of the longevity risk should be borne by those whose retirement incomes will be derived from the bonds.

Some might argue that the government should seek to charge a risk premium in excess of the economic cost. For example, if, in a Solvency II world, insurance companies writing annuity business end up having to hold capital in excess of true economic levels, because they are unable to hedge longevity risk, then they might be prepared to pay a premium price for longevity bonds if, by doing so, they can reduce their capital requirements. This would obviously depend on the Solvency II treatment of longevity bonds and the capital reduction that the regulators would allow.

It would be short sighted of governments to seek to exploit this arbitrage situation. If insurance companies can reduce their capital requirements closer to economic capital levels, then this should result in higher annuity values with the consequent benefits to government, pensioners and savers already highlighted.

In addition, we also believe that it is most unlikely that the market for longevity bonds will develop if the government just focuses on insurers. The bonds will need to be priced to attract DB pension plans which do not currently face solvency capital requirements. DB plans which do not have a pressing need for a full buy-out using annuities (which will be subject to Solvency II capital via insurers) and which want to engage in risk management will only buy longevity bonds if they believe they are priced fairly (and cheaper than longevity swaps and other derivative longevity hedges provided by the private sector). So, if we want to ensure DB pension plans buy longevity bonds issued by the government, the government should not price them above AAA.

Members in DC pension plans de-risking (i.e., life-styling or life-cycling) in the run up to their retirement also will have a choice between using long-dated bonds and longevity bonds and again many will be discouraged from using longevity bonds if the government looks to charge a mark-up beyond the fair price. Other investors, including investment banks, will also be discouraged from buying longevity bonds if they believe the longevity risk premium is excessive, because they will fear that the bonds will eventually fall in value to reflect their true economic cost.

So for the market in longevity bonds to take off, we believe they should be priced according to economic capital principles. The analysis below is intended to initiate the process of defining what is the fair economic price. Our intention is not to determine that price; rather it is to indicate one possible approach and the issues that need to be resolved for determining what the fair price might be. The approach we have adopted builds on the insurance industry ‘cost-of-capital’ method.³⁸ This determines a risk margin for capital above the best estimate of the value of the liabilities. The best estimate of the value of the liabilities in our model is derived from the median scenario and, at any point in time, is the present value of the expected future coupons on the bond from the median scenario discounted at the risk-free rate. The cost-of-capital method involves four stages:

- Determine the required credit rating for the bond.
- Project the longevity risk capital required for each year in the life of the bond to maintain the required credit rating.
- Multiply each annual capital requirement by a percentage cost of capital to give the cost of capital in money terms.
- Calculate the present value of each of these cost-of-capital amounts using a risk-free discount rate and sum to give the present value of the overall risk premium.

The starting point for quantifying the minimum risk premium that the government should charge to ensure intergenerational fairness is to consider the notional level of capital it would need to hold to achieve at least a AAA rating. It is important to realize that the government will not actually hold this capital – unlike an insurer – but simply uses the notional required capital amount to calculate the cost of capital for each year of the bond’s life. To calculate this notional capital, we ideally need to use stochastic mortality and interest rate modelling to determine the amount of notional capital that

³⁸ Chief Risk Officer Forum (2008). See Appendix C for an explanation.

would apply throughout the duration of the bond to ensure the bond's payments would be made with a continuing AAA level of confidence.

Our first task is to derive the survival probability on AAA bonds. We assume a yearly survival probability of 0.9995 in the analysis below to reflect the high standard of security that would be associated with government-issued longevity bonds. This is marginally higher than the annualized 20-year survival rates on AAA bonds of 0.9991 between 1970 and 2008 and 0.9994 between 1920 and 2008.³⁹

We then used the CBD model to project 10,000 longevity scenarios for English and Welsh males aged 65 at the end of 2006 (as shown in Figures 2, 3 and 5) and these were, in turn, used to calculate 10,000 present values of the coupon payments on a range of different types of longevity bond. Table 1 shows the distribution of life expectancies for males aged 65 and 75 at the end of 2006, according to the CBD model and quantiles of the distributions of longevity bond present values, payable immediately (PV(65,65) and PV(75,75)), payable from age 75 (PV(65,75)), payable from age 85 (PV(75,85)) and payable from age 90 (PV(65,90) and PV(75,90)), respectively.⁴⁰ For convenience, the median present value for each bond has been rescaled to £100 by adjusting the base coupon. A fixed risk-free discount rate of 4% is assumed throughout.⁴¹ Further, no allowance is made for expenses and other operational risks, since we are looking to quantify the pure price of the risk premium for longevity.

³⁹ The desired survival probability could be higher if required.

⁴⁰ Notice that the PV(65,90) bond is more volatile than the PV(65,75) bond which, in turn, is more volatile than the PV(65,65) bond. This is for precisely the same reason that a zero-coupon bond is more volatile than a coupon-paying bond with the same maturity: because the zero's cash flows are more heavily concentrated towards the end of its maturity than a bond paying regular coupons, it has greater duration.

⁴¹ The explanation for the choice of a fixed risk-free discount rate of 4% is given in Appendix C. A more sophisticated approach would stochastically model the risk-free term structure.

Table 1: Distribution of life expectancies and longevity bond present values

Quantile	e65	PV(65,65)	PV(65,75)	PV(65,90)	e75	PV(75,75)	PV(75,85)	PV(75,90)
0.005	18.77	94.68	88.02	60.36	10.96	93.28	79.06	66.04
0.01	18.93	95.22	89.14	63.55	11.07	93.94	81.34	69.40
0.025	19.17	95.97	90.81	68.42	11.20	94.81	83.82	73.22
0.05	19.37	96.57	92.19	72.44	11.34	95.67	86.48	77.63
0.5	20.51	100.00	100.00	100.00	12.03	100.00	100.00	100.00
0.95	21.82	103.65	108.39	134.43	12.79	104.57	114.76	126.10
0.975	22.07	104.34	109.98	141.43	12.94	105.37	117.62	131.67
0.99	22.38	105.12	111.73	150.07	13.14	106.57	121.17	138.73
0.995	22.57	105.63	113.03	155.36	13.28	107.31	123.87	143.24
Mean	20.53	100.03	100.09	101.25	12.04	100.05	100.19	100.65
Median annuity factor		12.619	5.222	0.675		8.420	2.106	0.815
Base coupon (£)		7.925	19.149	148.133		11.876	47.493	122.730

Notes: Derived from the CBD model estimated on English and Welsh male data for age 65 over the period 1991-2006. e65 and e75 = life expectancy at ages 65 and 75. PV(65,65) = present value of a bond with base coupon of £7.925 for a male aged 65, payable from age 65. PV(65,75) = present value of a bond with base coupon of £19.15 for a male aged 65, payable from age 75. PV(65,90) = present value of a bond with base coupon of £148.13 for a male aged 65, payable from age 90. The discount rate is assumed to be a risk free 4%. The median annuity factor is the present value of a base coupon of one unit payable yearly in arrears multiplied by the proportion of the cohort still alive at the end of each year, for the life of the annuitant from a given age. The base coupon is derived by dividing the median price of the bond (set as 100) by the median annuity factor. The actual coupon in each year a coupon is due is equal to the (rescaled) base coupon multiplied by the percentage of the population surviving between the bond's issue date and the coupon payment date.

We now need to determine the relevant quantiles of the distribution of present values to achieve a AAA rating. We do this at the undiscounted mean term of the expected payments.⁴² Table 2 shows the mean term on the issue date for a range of different

⁴² An alternative would have been to use the discounted mean term or duration of the bond. This, however, has the effect that it changes when the discount rate changes. This is inappropriate because the potential dispersion of projected cash flows, and hence the risk against which capital is being held, does not depend on interest rates. We did, however, examine the effect of using the discounted mean term with a fixed discount rate of 4% and it made very little difference to the final estimate of the longevity risk premium.

bonds. The corresponding AAA quantiles are shown in the last column. These are found by raising the survival probability of 0.9995 to the power of the mean term.

Table 2: Mean term of longevity bonds and corresponding AAA quantiles on issue date of bonds		
Longevity bond	Mean term	AAA quantile
LBM(65,65)	13.21	0.99341
LBM(65,75)	19.73	0.99018
LBM(65,90)	30.51	0.98486
LBM(75,75)	8.72	0.99565
LBM(75,85)	16.00	0.99203
LBM(75,90)	19.87	0.99011


Notes: The mean term is found by summing the expected coupons on a bond weighted by the number of years ahead each coupon occurs and then dividing by the sum of the expected coupons. The corresponding AAA quantile is found by raising the survival probability of 0.9995 to the power of the mean term. For example, for the LBM(65,65) bond, the mean term is 13.21 years and the corresponding AAA quantile is $0.9995^{13.21} = 0.99341$.

Using the information in Tables 1 and 2, we can determine the initial notional capital that is required for a AAA rating and then use this to calculate the cost of capital for each year of the bond's life.

Take, for example, the LBM(65,75) bond (i.e., one based on males age 65 with payments starting at age 75). On the issue date, the mean term is 19.73 years and therefore the AAA capital requirement can be derived from the 0.99018 quantile (see Table 2), giving an initial capital requirement of 11.73% (see Table 1 – the 0.99 quantile is £111.73, while the median is £100). Figure 6 shows graphically the level of economic capital required for the first year.

Figure 6: Distribution of 10,000 scenarios of the present values of 10-year Deferred Longevity Bond payments for males aged 65

Longevity Bond with coupon of £19.15 adjusted for survivorship of age 65 cohort



Note: See note to Figure 2

For subsequent years, we continue to use the best estimate of the bond’s coupons from the median scenario. However, we need to re-run the CBD model to produce new sets of 10,000 scenarios for each year in the future. In doing this, we assume that mortality rates follow the best estimate path from the median scenario up to the year (and associated age) that we are modelling and then we produce a new stochastic distribution of outcomes using drift and volatility parameters consistent with the CBD model used in the first year.

Although this results in a narrowing funnel of doubt as each year passes,⁴³ the mean term of the expected cash payments also reduces and this requires higher quantiles of the distribution to be used each year to maintain the desired AAA credit rating for the bond.⁴⁴ The net outcome of these opposing effects results in a lower capital mark-up percentage over time. Table 3 shows a subset of the mean terms, the resultant AAA quantiles and the capital mark-up percentages for LBM(65,75) and LBM(75,85) that can be applied to the series of best estimate liabilities derived from the median scenario.

⁴³ As the age 65 and 75 cohorts grow older, the range of possible outcomes narrows.

⁴⁴ This follows because 0.9995 raised to the power of a lower mean term produces a higher quantile than 0.9995 raised to the power of a higher mean term as Table 2 shows.

It is therefore possible using the CBD model to calculate the notional required AAA capital holdings for longevity risk for each year for any bond. We now need to multiply each one of these by the cost of capital and a risk-free discount factor and sum this series to produce the required risk premium which can be expressed as a percentage of the expected bond price of 100. We can then convert this to an effective basis points reduction from the risk-free rate.

Table 3: Mean term, AAA quantiles and resultant AAA capital as a percentage of best estimate liabilities						
Age	LBM(65,75)			LBM(75,85)		
	Mean term	Quantile	Capital %	Mean term	Quantile	Capital %
65	19.73	0.99018	11.73%			
70	14.73	0.99266	11.31%			
75	9.73	0.99515	11.01%	16.00	0.99203	21.81%
80	8.16	0.99593	10.34%	14.73	0.99266	20.70%
85	6.76	0.99663	10.05%	9.73	0.99515	19.89%
90	5.51	0.99725	9.66%	8.16	0.99593	18.31%
95	4.44	0.99778	9.04%	6.76	0.99663	17.05%
100	3.54	0.99823	8.52%	5.51	0.99725	15.82%
105	2.82	0.99859	8.07%	4.44	0.99778	13.98%
110	2.27	0.99887	7.57%	3.54	0.99823	12.90%

A critical factor in the process is to determine the appropriate cost of capital. This has been the subject of much debate in the run up to Solvency II: annuity companies are currently expected to use a 6% cost of capital when calculating their MVM. This is intended to cover a number of risk factors associated with annuity provision, the most significant being non-hedgeable longevity risk. However, the industry believes that this figure will lead to a SCR which will result in insurers being asked to hold capital above the true economic level.⁴⁵ The industry has therefore recommended a cost of capital in the range 2.5%-4.5% p.a., based on the cost of non-hedgeable risks and a capital level calibrated to a 0.995 survival probability over one year.⁴⁶ This

⁴⁵ Chief Risk Officer Forum (2008, pp. 16-18).

⁴⁶ Chief Risk Officer Forum (2008, p. 8). See Appendix C for an explanation.

approximately translates into a cost of capital in the range 1.67%-3% p.a., based on a 0.9995 annual survival probability.⁴⁷

The upper end of this range is substantially higher than a government would be expected to charge. This is because the longevity risk faced by governments is lower than that faced by insurers because they have the benefit of having a more reliable estimate of current longevity exposures. They therefore have a more accurate starting point for modelling longevity improvement risk. They also face less random variability in trend improvements in longevity as government-issued longevity bonds will be based on national population data. By contrast, the population relevant for insurers is a small and much more volatile subset of the national population. A case could therefore potentially be made for government to use a cost of capital of around 2%.^{48, 49}

Table 4 shows the total risk premium for a number of longevity bonds for illustrative costs of capital of 2% and 3%. It also shows the corresponding basis points reductions from the risk-free rate. Take LBM(65,75) and a 2% cost of capital, for example. This bond has a total risk premium of 3.2%. This means that the issue price of the bond would be £103.20. The effective yield on the bond is equal to the risk-free rate less the basis points reduction, so the effective yield on LBM(65,75) is 3.821%.⁵⁰

⁴⁷ Chief Risk Officer Forum (2008, Figure1, p. 30).

⁴⁸ This would include an allowance for model risk, e.g., in the model used to project future mortality rates.

⁴⁹ An alternative approach to the cost-of-capital method used in this paper is the 'percentile method' which determines the level of capital needed to ensure that all payments can be met for a set percentage of all the scenarios. In the context of Solvency II, a probability of 75% has been suggested. By using the initial 10,000 present value scenarios from the CBM model, a 75 percentile risk premium can be determined and, in turn, an implied cost of capital can be calculated. In this case, the percentile method implies costs of capital of 2.11% for LBM(65,75), 1.75% for LBM(65,90) 2.77% for LBM(75,85) and 2.45% for LBM(75,90).

⁵⁰ By using a discount rate of 3.821%, the present value of the coupon payments on the LBM(65,75) bond equals £103.20.

Table 4: Risk premiums and basis points reduction in yield on longevity bonds				
Bond	2% cost of capital		3% cost of capital	
	Risk premium	Bps reduction in yield	Risk premium	Bps reduction in yield
LBM(65,65)	1.4%	13.4 bps	2.0%	20.0 bps
LBM(65,75)	3.2%	17.9 bps	4.7%	26.5 bps
LBM(65,90)	15.1%	48.7 bps	22.6%	70.8 bps
LBM(75,75)	1.2%	16.5 bps	1.8%	24.7 bps
LBM(75,85)	4.1%	27.6 bps	6.2%	40.8 bps
LBM(75,90)	8.2%	42.6 bps	12.4%	62.2 bps

Notes: The risk premium is the total for each bond. The basis points reduction shows the annual reduction from the assumed risk-free yield of 4%.

VIII. Who benefits from government issuing longevity bonds?

Who benefits from governments assisting in encouraging the optimal sharing of longevity risk? The simple answer is everyone. Everyone should benefit from having a market price for longevity risk and the ability to hedge systematic longevity risk. But there are also more specific benefits.

The government:

- Gains by having both a more secure DC pension savings market and a more efficient annuity market, resulting in less means-tested benefits and a higher tax take.
- Should gain access to a new source of long-term funding which, by widening the investor base, lowers the cost of government issuance.
- Is able to issue bonds with a deferred payment structure to help its current funding programme and improve its cash flow.
- Earns a market-determined longevity risk premium thereby further reducing the expected cost of the long-term national debt.

For DB pension plans:

- Have the opportunity to reduce longevity risks.
- Can hedge longevity risk exposure prior to buy out.

Insurers:

- Can potentially establish a mark-to-market longevity risk term structure and hence hold the optimal level of economic capital or at least hold capital closer to the economic level.
- Longevity bonds will help insurers to play an aggregating role in providing pension plans and individuals with longevity insurance, whilst being able to pass on a proportion of their risk to the capital market; this would reduce their longevity concentration risk and facilitate the spread of longevity risk around the capital markets.

The capital markets:

- Get help to kick start market participation through the establishment of reliable longevity indices and key price points on the longevity risk term structure.
- Can build on this longevity risk term structure with liquid longevity derivatives.

Investors:

- Get access to a new (longevity-linked) asset class whose returns are uncorrelated with traditional asset classes, such as bonds, equities and real estate.

Regulators:

- A longevity risk term structure should help the insurers' regulator (the Prudential Regulation Authority⁵¹ in the UK) validate insurers' economic capital, thereby making regulation more robust.

⁵¹ This replaced the Financial Services Authority in April 2013.

- Longevity bonds should help an orderly transfer of longevity risk from DB plans to the capital markets, thereby reducing reliance on an uncertain sponsor covenant and reducing concentration risk amongst insurers, and, in turn, giving comfort to the pension plans' regulator.
- A longevity risk term structure should help facilitate the calculation of any risk-based levy to a pension insurance plan (the Pension Protection Fund in the UK).⁵²

Pension plan members:

- DB pension plan members potentially get better security.
- DC pension plan members get better valued annuities which produce a higher lifetime income when they retire.
- Further, individuals with DC pension plans would have a means of hedging the longevity risk associated with purchasing an annuity at retirement.

IX. Growing support for government issuance of longevity bonds

Support for governments to issue longevity bonds is growing steadily, not only in the UK, where the situation is most immediate, but also internationally.

The UK Pensions Commission suggested the government should consider the use of longevity bonds to absorb tail risk for those over 90 or 95, provided it exits from other forms of longevity risk pre-retirement which it has done by linking state pension age to increases in life expectancy and by raising the future state pension age from 65 to 68 by 2046. "One possible limited role for government may, however, be worth consideration: the absorption of the 'extreme tail' of longevity risk post-retirement, i.e., uncertainty about the mortality experience of the minority of people who live to very old ages, say, beyond 90 or beyond 95."⁵³

⁵² The Pensions Regulator in the UK is responsible for the regulation of occupational trust-based DB and DC schemes and attempts to limit the number of DB schemes needing support from the Pension Protection Fund (which was based on the US Pension Benefit Guaranty Corporation).

⁵³ Pension Commission (2005, p. 229).

The UK Confederation of British Industry (CBI), which represents British employers, has argued: “Government should drive development of a market in longevity bonds, a similar instrument to annuities, by which the payments on the bonds depend on the proportion of a reference population that is still surviving at the date of payment of each coupon. This should be done through limited seed capital and supporting policy work on the topic. Government could also consider how best to match government bond issues to pension scheme needs, including the provision of more long-dated bonds and whether government should issue mortality bonds itself.”⁵⁴

According to the OECD: “Governments could improve the market for annuities by issuing longevity indexed bonds and by producing a longevity index.”⁵⁵

The World Economic Forum has argued: “Given the ongoing shift towards defined contribution pension arrangements, there will be a growing need for annuities to enhance the security of retirement income. Longevity-indexed bonds and markets for hedging longevity risk would therefore play a critical role in ensuring an adequate provision of annuities.”⁵⁶

Finally, the IMF states: “Although the private sector will further develop market-based transfer mechanisms for longevity risk if it recognizes the benefits of doing so, the government has a potential role in supporting this market. Measures could include provision of better longevity data, better regulation and supervision, and education to promote awareness of longevity risk. Those governments that are able to limit their own longevity risk could consider issuing a limited quantity of longevity bonds to jumpstart the market.”⁵⁷

⁵⁴ *Redressing the Balance - Boosting the Economy and Protecting Pensions*, CBI Brief, May 2009.

⁵⁵ Antolin and Blommestein (2007).

⁵⁶ World Economic Forum (2009).

⁵⁷ International Monetary Fund (2012).

X. Counter arguments

While we feel we have put forward a number of strong arguments supporting the case for longevity bonds that are issued by governments, we do need to acknowledge and then address a number of counter arguments.

First, concerns have been raised that governments are not natural issuers of longevity bonds because of their large existing exposure – in excess of £5trn in the case of the UK government – to longevity risk.

We would argue that a government's exposure to unanticipated longevity improvements through the issuance of longevity bonds is – or at least could be – well hedged. First, the government receives a longevity risk premium from issuing the bonds. Second, in the event that the risk premium proves to be insufficient, the government can reduce its state pension spend and increase its pre-retirement tax take by raising the state pension age, as recommended by the UK Pensions Commission. The next generation might have to work longer, but will, in any case, have ended up being a fitter generation than anticipated and so be able to earn more income which, in turn, will produce more tax. Third, since the issuance of longevity bonds should result in a more efficient annuity market and hence higher incomes in retirement, this should also result in an increase in the tax take and help to reduce the amount of means-tested benefits. In addition, it should be noted that the higher tax take and lower means-tested benefits arising from a more efficient annuity market applies to the lifetimes of all pensioners buying an annuity, whereas the tail risk protection provided by deferred tail longevity bonds applies only to those surviving over 90, some 25 years in the future.

Overall, once a government is only issuing deferred tail longevity bonds, the risk will be very manageable and consistent with the government's role of facilitating intergenerational risk sharing. We believe that there could be a significant cost-benefit to the government from the issuance of longevity bonds and therefore a strong, indeed overwhelming, case for a government to issue longevity bonds.

The second criticism is that there is no role at all for a government in issuing longevity bonds as argued by Dowd (2003) and Brown and Orszag (2006).

Dowd (2003) criticized the original argument used by Blake and Burrows (2001) to justify government issuance of longevity bonds (or what Blake and Burrows called survivor bonds), namely the appeal to the Arrow-Lind Theorem on social risk bearing. This theorem states that by dispersing an aggregate risk across the population (of taxpayers) as a whole, the associated risk premium on a longevity bond issued by the government would be lower than that charged by a private-sector issuer. Dowd countered that many of the assumptions underlying the theorem – such as taxes are costless to collect, each household bears an equal share of the tax burden, and an absence of distributional effects – do not hold in practice. Instead, he argued that capital markets are better suited than any government to bear and share risks, since they allow risks to be diversified internationally. In short, Dowd argued that government intervention was unnecessary, since private-sector parties were perfectly capable of creating and trading longevity-linked instruments and derivatives themselves. There was no market failure for the government to correct, rather the time is not yet ripe: “The fact that a particular innovation has not yet occurred does not in itself constitute an argument for government intervention to bring it about. Any good new idea, including that of survivor derivatives, should eventually take off – but we have to give it time.... When the time is ripe, it is therefore entirely possible, and even likely, that markets for survivor derivatives – survivor bonds, forwards, futures, options and swaps, and annuity securitization – will take off, and eventually become as familiar as comparable instruments such as credit derivatives are today” (pp. 347-8).

Brown and Orszag (2006) also accept that a longevity risk premium would need to be paid in order to hedge aggregate longevity risk, but they argue that it is not sufficiently high to cause a market failure and hence justify government intervention: “we suspect that this risk does exert some upward pressure on annuity pricing, possibly in the range of a few percentage points” (p. 622). They also accept that the intergenerational sharing of longevity risk can potentially improve social welfare. Suppose a scientific discovery improves the life expectancy of all current and future generations. Current 80-year olds would be unable to respond to this by re-entering

the labour market and hence would experience a lower standard of living as their remaining wealth would have to be spread over a longer period. Younger generations are more able to adjust to this mortality shock. Hence the financial risk from such a shock could be spread over a number of generations and this would improve social welfare. Since only the government is able to enforce intergenerational contracts, there is a potential role for the government in efficiently spreading risk across generations. However, Brown and Orszag believe that it is unlikely that the government will spread risk efficiently: “to maximize social welfare, it is not sufficient that the government move *any* amount of risk from the current generation to some other generation. Rather, the government needs to move the *optimal* amount of risk onto the right generations” (p. 625). Instead, they believe that the government will favour the current generation of voters, and particularly the large number of vocal grey voters, over generations as yet unborn, by transferring “more than the optimal amount of risk to future generations” (p. 629).⁵⁸

We would argue that there is a role for both government and the private sector in developing a longevity market. As discussed in Figure 1, the private sector is best at hedging specific longevity risk, once it has hedged systematic longevity risk. The government is the only agent in society with both the capacity and credibility to provide a long-term hedge for systematic longevity risk through the issuance of longevity bonds. While Dowd, Brown and Orszag highlight some of the difficulties associated with the government’s ability to forecast future mortality improvements, the existence of longevity bonds would provide an incentive for the government to collect better death records and improve its longevity forecasting techniques, both of which would have wider social benefits. Even if the private sector is better at forecasting than the government – which in this case is hard to believe since it is the government that collects death statistics – systematic longevity is a slowly building trend risk and the private-sector issuer of a longevity bond risks insolvency if it gets that trend wrong in a way that the government will its unlimited powers of taxation does not.

⁵⁸ Dowd (2003, pp. 346-7) makes the same point: “The intergenerational argument is open to the objection that governments have an incentive to put the interests of current voters ahead of those of future voters”. We would argue that the issuance of longevity bonds would help to reduce this incentive. The current generation is getting its longevity risk insurance for free: if longevity bonds were issued, it would have to pay for it!

The third criticism is that even if longevity bonds are issued by the government, there is a question mark concerning the potential liquidity of the market trading longevity bonds. Some have argued that liquidity is likely to be thin, since any new information concerning mortality that would be sufficiently significant to motivate trading is likely to arrive very infrequently. While this is true, we believe that there are important lessons from the inflation-linked financial futures market. Early attempts to introduce such a market were initially unsuccessful but they eventually succeeded and inflation indices have similar characteristics to longevity indices, especially in their low frequency of publication.

The first attempt occurred when CPI futures contracts were listed on the US Coffee, Sugar and Cocoa Exchange in June 1985. This contract was delisted in April 1987, with only 10,000 contracts ever having been traded. The key reasons for the failure of this contract were: there was no underlying inflation-linked securities market at the time, the underlying was an infrequently published (i.e., monthly) index, and there was no stable pricing relationship with other instruments to attract the attention of arbitrageurs. The second attempt occurred when Treasury inflation-protected securities (TIPS) futures were 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. The key reasons for the failure of this contract were: TIPS had only started trading five months before, there was just a single 10-year TIPS trading, the futures contract competed with the underlying for liquidity, and there was uncertainty over the future of the TIPS program. The final attempt was in February 2004 when the Chicago Mercantile Exchange launched a CPI futures contract which is still trading. The reasons for the success of this contract are: inflation-linked securities have gained acceptance amongst investors, TIPS have evolved into recognized asset class, there is a well-understood pricing relationship allowing for arbitrage opportunities between TIPS, fixed-interest Treasury bonds and CPI futures, the US Treasury is committed to long-term TIPS issuance, CPI futures do not compete directly with but rather complement TIPS and use same the inflation index, and liquidity is enhanced by electronic trading on Globex. This experience therefore suggests that it is possible to create a liquid market in an instrument based on an infrequently published index.

The fourth criticism is that longevity bonds are unnecessary since the load in annuity prices is sufficiently large to a) absorb the increase in regulatory capital that will be required after the introduction of Solvency II in the absence of longevity bonds, and b) to absorb the longevity risk in countries not subject to Solvency II (e.g., the US and Australia).

Our response is that there is limited scope for annuity providers to absorb either the costs of the additional capital requirements or the aggregate longevity risk without seriously reducing the money's worth of the annuities they sell.⁵⁹

The life annuity market in the UK has scale (a £12bn per annum market - around a half of the global annuity market) and as a consequence is price competitive with a number of life insurers competing for business. It is relatively easy for pensioners to compare the different guaranteed incomes on offer in exchange for their pension savings.

In recent years, the money's worth of the UK annuity market has been assessed and tracked by Professors Edmund Cannon and Ian Tonks. They were commissioned by the Department for Work and Pensions (DWP) in 2009 to produce a detailed report on the money's worth of annuities in the UK. Their report examines a time series of pension annuity rates in the UK for the period 1994 to 2007. "The report computes the money's worth of annuities and finds that, on average, the money's worth over the sample period for 65-year old males has been 90 per cent, and for 65-year old females has been a similar but slightly larger 91 per cent. Taking into account load factors associated with annuity contracts and in comparison with other financial and insurance products this implies that annuities are fairly priced." (Cannon and Tonks (2009, xiii).

⁵⁹ The conventional methodology for valuing annuities is to calculate the 'money's worth' statistic, which will equal 100% when annuity providers have no administrative costs and are making no profits. In practice, the money's worth is typically less than 100 per cent due to the presence of administrative costs, risk charges (in form of cost of capital) and the need for annuity providers to make a 'normal profit'. The sum of the costs and normal profit is called the 'load factor'.

Cannon and Tonks' analysis shows that there is some evidence that the money's worth has fallen since 2002. They discuss a number of reasons for this, including: changes in insurance regulation, changes in industrial concentration, an insurance cycle, the pricing of mortality uncertainty, and the growth in the impaired lives market. The last of these is becoming an increasingly important factor in the UK and it has resulted in the money's worth for standard annuities (i.e., those for healthy lives) falling as insurance companies have made allowance for the selection effects caused by the introduction of enhanced rates for pensioners with health impairments that reduce their expected life expectancy. Around 30% of pensioners qualify for enhanced annuity rates and life insurers have adjusted the rates on standard annuities to reflect the longer life expectancy of the 70% buying standard annuities. The other main reason is that UK insurers have increased the loading for the cost of their risk capital to reflect the fact that they expect to have to hold more capital in a Solvency II world. This trend has accelerated since 2009 as the introduction of Solvency II comes nearer. In short, the load in annuities cannot take much more strain without adversely impacting the size of the annuity payments.

The fifth and final criticism that we consider is that basis risk is sufficiently large that it would negate any gains from holding longevity bonds.

We recognise that basis risk is an important issue. There will be a requirement under Solvency II for annuity companies to hold capital to cover basis risk where they have a hedging instrument that is not perfect. However, given that no longevity bonds have yet been issued, no annuity provider has been in a position to agree the scale of capital required with its regulator. The level of capital will clearly depend on the composition and size of the insurer's annuity population. However, reinsurers who are also caught by Solvency II would be more able to consolidate exposure by pooling portfolios from different providers and therefore suffer less basis risk. It is possible that reinsurers could end up using longevity bonds to manage their longevity risk and reduce their Solvency II capital requirement, whilst providing indemnity rather than indexed solutions to insurers with small pools of annuities.

Whilst it is hard to be absolutely sure at this stage in the development of the market, we do not believe that basis risk means that longevity bonds will be ineffective. Basis

risk arises in other markets where imperfect hedging instruments are used, such as interest rate and currency futures contracts. Using these contracts leads to both contemporaneous and time basis risk, but this does not prevent them from providing highly effective – if not perfect – hedges as was discussed in detail in Blake et al. (2006).

XI. Next Step

If we accept that longevity bonds have a potentially important role to play in hedging systematic longevity risk, then the next step is for governments in countries with significant private sector pension funds to set up a working party to undertake a cost-benefit analysis of government issuance of longevity bonds to help manage the associated longevity risk exposure. The terms of reference of this working party should cover the benefits that would accrue, the scale of the longevity risk that governments would be assuming, and the actions governments can take to mitigate this risk. The working party should also work through the practicalities of issuing longevity bonds, including the construction of reference longevity indices, potential demand, pricing, liquidity and taxation.⁶⁰

Appendix A: A Brief Guide to Solvency II

Solvency II is similar to the banks' regulatory regime Basel II, and its purpose is to align regulatory capital more closely with economic capital. It is due to come into force in all member states of the European Union in 2014, having already been delayed several times.

The European Commission's Solvency II initiative to improve the regulation of European insurance companies started in 2000. Its aim is to ensure improved risk management and greater consistency in the calculation of capital requirements across European insurers.

⁶⁰ Longevity bonds are annuity bonds with the coupon payment involving a return of capital element as well as an interest element. The tax treatment will therefore be more complicated than with a conventional bond.

The capital that needs to be held under the current Solvency I framework is calculated using simple formulae which result in different levels of prudence for different insurance products and even different portfolios within product categories. As a result, it is difficult to compare the financial strength of European insurers.

Under Solvency II, Insurers will be required to hold a minimum Solvency Capital Requirement (SCR) which is calculated to ensure that the firm holds sufficient capital to cover against adverse events occurring over the next year with a probability of 99.5%.

The use of a one-year value-at-risk measure reflects a desire by EU regulators for consistency with the Basel capital adequacy regime for banks, although many insurance experts would argue that this is flawed given the long-term and different nature of insurance liabilities.

Insurance firms can either use Standard Formulae or develop their own Internal Models to calculate their SCR. Both methodologies require the firms to use assumptions set by the EU regulator regarding the valuations of assets and liabilities. To ensure consistency and maximum harmonisation across EU member states, national regulators will have the responsibility to ensure that their insurers use the final EU-wide standardised assumptions and methodologies.

The objective of the Solvency II valuation approach is to enhance comparability and transparency across European insurers. The Committee of European Insurance and Occupational Pension Supervisors (CEIOPS) has been advising on the development of common Solvency II risk margin calculation methodology and assumptions. This has proved to be a difficult and contentious task given the diversity of products and current practices across member states and there are still a number of unanswered issues particularly concerning annuity business.

Where possible a mark-to-market approach is used. However, if there is no deep and liquid financial market resulting in risks that are non-hedgeable then a mark-to-model approach is used. Longevity risk is currently deemed to be non-hedgeable.

The calculation of the risk margin for a non-hedgeable risk is based on the cost-of-capital (CoC) method, with CoC defined as the cost of holding sufficient capital consistent with projected future SCRs to support the business. Under the CoC approach, the CoC charge in every period should be calculated by multiplying the projected capital requirement in respect of non-hedgeable risk capital by a predefined CoC rate. This is the philosophy we have attempted to mirror in calculating the longevity bond prices in Section VII. However as this CoC approach requires complex multi-year risk modelling, it is expected that some simplification will be allowed under Solvency II. The proposed Solvency II CoC of 6% above the risk free rate has also been challenged by the Chief Risk Officers' Forum⁶¹ and others.

A firm date for the introduction of Solvency II has still not been finally fixed and there are still a number of uncertainties particularly for annuity providers around the allowance for illiquidity premiums and future longevity risks.

Finally, it is important to reiterate that our proposal for governments to issue longevity bonds is not primarily a response to Solvency II in the EU. Our key argument is that longevity risk is an inter-generational risk that requires governments in all countries to help to manage.

Appendix B: The Cairns-Blake-Dowd Model

The Cairns-Blake-Dowd (CBD) (2006) model is a two-parameter stochastic mortality model that fits the logit of the mortality rate to the two factors as follows:

$$\text{logit}(q(t, x)) = \log \left[\frac{q(t, x)}{1 - q(t, x)} \right] = \beta_x^{(1)} \kappa_t^{(1)} + \beta_x^{(2)} \kappa_t^{(2)}$$

where $q(t, x)$ is the mortality rate at time t and at age x , $\kappa_t^{(i)}$ is the i^{th} time-varying factor that drives the dynamics of mortality rates, and $\beta_x^{(i)}$ is the i^{th} age-related weight on $\kappa_t^{(i)}$. The CBD model adopts very simple parametric forms for the age-related weights:

⁶¹ See Appendix C for further information on the C-o-C Method

$$\beta_x^{(1)} = 1$$

$$\beta_x^{(2)} = (x - \bar{x})$$

where $\bar{x} = n^{-1} \sum_i x_i$ is the mean age in the sample range and n is the length of the sample range. This particular parameterization means that the first time-varying factor influences the level of the mortality term structure at time t , while the second influences the slope.

A number of studies have shown that the CBD model fits mortality rate data well at high ages (above 50) in terms of goodness-of-fit, backtesting and the generation of mortality density forecasts (see, e.g., Cairns et al. (2009, 2011) and Dowd et al. (2010a,b)).

Appendix C: The Cost-of-Capital Method and a Justification for the Cost-of-Capital Assumptions used to Price the Longevity Bond

Our model for pricing longevity bonds makes use of the ‘cost-of-capital’ method outlined in the Chief Risk Officer (CRO) Forum’s (2008) report ‘Market Value of Liabilities for Insurance Firms – Implementing Elements for Solvency II’. This report addressed both core principles and practical issues relating to the calculation of the market value of liabilities under Solvency II.

By the ‘cost of capital’ (CoC), we mean the cost above the risk free rate. As shown in Table 4, the CoC can be expressed as a risk premium above or as a reduction in yield from the risk free rate. We can interpret the CoC as the longevity risk premium demanded by government to ensure inter-generational fairness, as discussed in Section V.C.

The CRO Forum sought advice from Dr Philipp Keller of Ernst & Young and Professors Shaun Wang and Richard Phillips of Georgia State University concerning the calibration of the CoC. The resulting 2008 report concluded (pages 8 and 18): “Research commissioned by the CRO Forum suggests that a suitable range for the

cost of capital rate is 2.5% - 4.5% per annum. This rate is intended to be applied to an Solvency Capital Requirement (SCR) calibrated to a 99.5% confidence interval over a one year time horizon.” Figure 1 on page 30 of the report shows the CoC rate as a function of confidence level in the base case that they assumed: “It can be seen that the CoC rate reduces as the level of capitalisation increases, reaching a level of COC (99.99%) = 2.6% for AAA-rated companies.”

The CRO Forum’s base case also assumed a risk free rate of 4%, hence our use of this rate in our study. Figure 6 in the CRO report on page 35 shows the sensitivity of the cost of capital as a function of the confidence level for a range of risk free rates. An 8% risk free rate suggests a 3.5% CoC, a 5% risk free rate a 2.5% CoC, and a 2% risk free rate a 2% CoC, all at the 99.99% one year confidence level.

The CRO Forum’s analysis of and charts on The CoC lend support for our decision to show the longevity bond pricing at COCs of 2% and 3%, particularly when we are calculating capital at the 99.95% one year confidence level. The quantum of economic capital at this level is much higher than at the 99.5% level which is consistent with the use of a lower cost of capital.

References

Antolin, Pablo, and Hans Blommestein (2007), *Governments and the Market for Longevity-Indexed Bonds*, OECD Working Papers on Insurance and Private Pensions, No. 4, OECD Publishing, Paris.

Blake, David, and William Burrows (2001), “Survivor Bonds: Helping to Hedge Mortality Risk”, *Journal of Risk and Insurance*, 68: 339-348.

Blake, David, Andrew Cairns and Kevin Dowd (2006), “Living with Mortality: Longevity Bonds and other Mortality-Linked Securities”, *British Actuarial Journal*, 12: 153-228.

Blake, David, Andrew Cairns and Kevin Dowd (2008), “Longevity Risk and the Grim Reaper’s Toxic Tail: The Survivor Fan Charts”, *Insurance: Mathematics & Economics*, 42: 1062-66.

Bohn, Henning (2012), “Private versus Public Risk Sharing: Should Governments Provide Reinsurance?”, in A. Lans Bovenberg, Casper van Ewijk and Ed W.M.T. Westerhout (2012) *The Future of Multi-Pillar Pensions*, Cambridge University Press, Cambridge.

Brown, Jeffrey, and Peter Orszag (2006), “The Political Economy of Government-Issued Longevity Bonds”, *Journal of Risk and Insurance*, 73: 611-631.

Cairns, Andrew, David Blake and Kevin Dowd (2006), “A Two-Factor Model for Stochastic Mortality with Parameter Uncertainty: Theory and Calibration”, *Journal of Risk and Insurance*, 73: 687-718.

Cairns, Andrew, David Blake, Kevin Dowd, Guy Coughlan, David Epstein, Alen Ong, and Igor Balevich (2009), “A Quantitative Comparison of Stochastic Mortality Models Using Data From England and Wales and the United States”, *North American Actuarial Journal*, 13(1): 1-35.

Cairns, Andrew, David Blake, Kevin Dowd, Guy Coughlan, David Epstein, and Marwa Khalaf-Allah (2011), “Mortality Density Forecasts: An Analysis of Six Stochastic Mortality Models”, *Insurance: Mathematics and Economics*, 48: 355–367.

Cannon, Edmund, and Ian Tonks (2009) “Money’s Worth of Pension Annuities”, Department for Work and Pensions, Research Report No 563.

Coughlan, Guy, Marwa Khalaf-Allah, Yijing Ye, Sumit Kumar, Andrew Cairns, David Blake and Kevin Dowd (2011), “Longevity Hedging 101: A Framework for Longevity Basis Risk Analysis and Hedge Effectiveness”, *North American Actuarial Journal*, 15(2): 150-176.

Chief Risk Officer Forum (2008), “Market Value of Liabilities for Insurance Firms: Implementing Elements for Solvency II”, July. [www.thecroforum.org/publication/market_value]

Dowd, Kevin (2003), “Survivor Bonds: A Comment on Blake and Burrows”, *Journal of Risk and Insurance*, 70: 339-348.

Dowd, Kevin, David Blake, Andrew Cairns and Paul Dawson (2006), “Survivor Swaps”, *Journal of Risk and Insurance*, 73: 1-17.

Dowd, K., Andrew Cairns, David Blake, Guy Coughlan, David Epstein, and Marwa Khalaf-Allah (2010a), “Evaluating the Goodness of Fit of Stochastic Mortality Models”, *Insurance: Mathematics and Economics*, 47: 255–265.

Dowd, K., Andrew Cairns, David Blake, Guy Coughlan, David Epstein, and Marwa Khalaf-Allah (2010b), “Backtesting Stochastic Mortality Models: An Ex-Post Evaluation of Multiperiod-Ahead Density Forecasts”, *North American Actuarial Journal*, 14 (3): 281-298.

Hobbs, David (2012), “A Broader Picture of the Public Sector Balance Sheet: State Pension and Other Pension Obligations - An Update at April 2012”, Office for National Statistics, 27 April (www.ons.gov.uk/ons/dcp171766_263808.pdf)

International Monetary Fund (2012), *The Financial Impact of Longevity Risk*, Chapter 4 of *Global Financial Stability Report*, April, Washington DC.

Levy, Sarah (2012) “Pensions in the National Accounts – A Fuller Picture of the UK’s Funded and Unfunded Pension Obligations”, Office for National Statistics, 27 April (www.ons.gov.uk/ons/rel/pensions/pensions-in-the-national-accounts/uk-national-accounts-supplementary-table-on-pensions--2010-/art-mainarticle.html)

Milevsky, M., S. Promislow, and V. Young (2006), “Killing the Law of Large Numbers: Mortality Risk Premiums and the Sharpe Ratio”, *Journal of Risk and Insurance*, 73: 673-686.

OECD (2011) *Pension Markets in Focus*, Issue 8, July, Paris.

Pension Commission (2005) *A New Pension Settlement for the Twenty-First Century*, HMSO: Norwich.

Pension Protection Fund and the Pensions Regulator (2006), *The Purple Book: DB Pensions Universe Risk Profile 2006*, Croydon and Brighton.

Richards, Stephen J., J. G. Kirkby, and Ian. D. Currie (2006), “The Importance of Year of Birth in Two-Dimensional Mortality Data”, *British Actuarial Journal*, 12: 5–38.

Tully, Andrew (2011) “How Solvency II will affect the Annuity Market”, *Retirement Planner*, 22 June.

Willetts, Richard C. (2004), “The Cohort Effect: Insights and Explanations”, *British Actuarial Journal*, 10: 833–77.

World Economic Forum (2009) *Financing Demographic Shifts Project*, June, Geneva.