DISCUSSION PAPER PI-9910

Survivor Bonds: Helping to Hedge Mortality Risk

David Blake and William Burrows

October 2001

ISSN 1367-580X

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

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

©*The Journal of Risk and Insurance*, 2001, Vol. 68, No. 2, 339-348

Survivor Bonds: Helping To Hedge Mortality Risk

David Blake
William Burrows

Abstract

One of the key problems facing annuity providers is mortality risk, the risk of underestimating mortality improvements. The authors argue that the government could help the issuers of annuities to hedge aggregate mortality risk by introducing a new type of bond, which the authors call a survivor bond. The future coupons on this bond depend on the percentage of the population of retirement age on the issue date who are still alive on the future coupon payment dates. The coupons on the bond therefore decline over time but continue in payment until the last members of this population cohort have died. The government would therefore be assuming a risk that has hitherto been borne by the private sector. However, governments now issue inflation-indexed bonds, and the authors would argue that inflation risk is a much greater risk than mortality risk in aggregate. Furthermore, governments directly contribute to mortality risk: for example, they promote public health campaigns that, if successful, lead to mortality improvements that are difficult to predict many years ahead. Survivor bonds enable pension provision to be a shared responsibility between the public and private sectors.

Introduction

Governments throughout the world are reviewing their pension systems with a view to improving the arrangements for those who are poorly served by existing systems. However, the debate to date has focused almost exclusively on the *accumulation stage*, the period of active membership of a pension plan up until the retirement date. Issues here include the advantages and disadvantages of funding versus pay-as-you-go (PAYG), defined benefit versus defined contribution, and active versus passive fund management.[1]

David Blake is professor of financial economics and director of the Pensions Institute, Birkbeck College, University of London. William Burrows is business development director, Prudential Annuities. The authors wish to thank (without implicating) Patrick Brockett (the editor), Angus MacDonald, Olivia Mitchell, Mike Orszag, Tim Sheldon, Guy Thomas, Steve Zeldes, and four anonymous referees for useful insights and comments.

[1] For a review of issues involved in the debate in the U.K., see Blake and Orszag (1997) or Blake and Orszag (1998).

What has received much less attention is the *retirement stage*. This lack of focus on retirement issues should not be surprising. When a new plan is established, the retirement stage is typically 40 years or more ahead. Furthermore, for members of defined benefit plans (whether public sector or private sector, funded or unfunded), someone else apart from the plan member is guaranteeing or at least promising to deliver a particular level of pension in retirement. However, this is not the case with defined contribution (DC) plans, and most new plans being established throughout the world are DC plans. With DC plans, there is no guaranteed pension at retirement. The retiree must live off whatever fund value has been accumulated at the time of retirement. Furthermore, because of uncertain life expectancy, individuals face the risk of outliving their resources. They can insure against the risk of living too long by buying an annuity from an insurance company. The purchase of a pension annuity is mandatory in some countries such as the U.K., where a DC plan member must purchase an annuity at some point between the ages of 50 and 75. It is not compulsory in other countries, such as the U.S., Australia, and Germany.

However, there is a key problem that annuity providers must face. Mortality has improved substantially this century, and it is very difficult to forecast improvements in mortality accurately. This is what is meant by *mortality risk*. Under a state-run PAYG system, the government bears this mortality risk. It is one thing for a government or a large occupational plan with the ability to adjust contribution rates to bear such a risk; it is quite another to expect an insurance company, particularly a new insurance company being established in a country with no previous history of annuity provision, to accept this risk unconditionally. There are other hidden costs: e.g., because lifetimes are uncertain, insurance companies must construct hedged investment portfolios consisting of many types of long-term bonds. Such portfolios are costly to manage and, in any case, provide imperfect hedges of mortality risk.

This article investigates how mortality risk can be managed. The authors propose that the government issue a new type of a bond, a *survivor* (or *life annuity*) *bond*, which allows its holders to hedge aggregate mortality risk and so reduce the management charges associated with constructing a hedged investment portfolio.

The next section analyzes mortality risk in more detail. The third section presents the authors' policy proposal, while the fourth section examines the implications of introducing survivor bonds.

MORTALITY RISK

Consider an individual who will live for exactly T additional years. This individual could use a lump sum to purchase an annuity from an insurance company or buy an annuity bond (which pays coupons only and has no principal repayment) directly from the financial markets; both will yield a constant income stream for T periods. In the absence of arbitrage, both investments should cost the same. The market price of a T-year annuity bond is (Blake, 2000):

$$P = \frac{d}{r}\left[1 - (1 + r)^{-T}\right], \tag{1}$$

where d is the annual coupon and r is the relevant discount rate (as a proportion). If someone purchases this bond at price P and then lives exactly T years, this is equivalent to someone purchasing a T-year annuity for an amount P that pays d per year in arrears.

In reality, neither individuals nor insurance companies know exactly how long an individual annuitant will live. In the case of an annuity bond that continues to pay out coupons for as long as the individual is alive, its price depends on the whole probability distribution of death rates for this individual; in other words, T is a random variable and not a fixed parameter. As a consequence, the market price of such annuity bonds depends on expectations about the random variable T:

$$P = E\left[\frac{d}{r}\left[1 - (1 + r)^{-T}\right]\right], \tag{2}$$

where E is the expectations operator.

Annuity bonds with random maturities are currently not traded on financial markets, but insurance companies still provide life annuities with uncertain values of T. Each insurance company will attempt to minimize its exposure to mortality risk by holding a portfolio of fixed-term bonds that matches the anticipated mortality profile of its annuitants and by building up a large enough pool of annuitants to reduce the probability that payouts will be worse than expected.

However, an insurance company cannot predict mortality perfectly. To consider the effects of errors in forecasting mortality improvements, the authors denote the probability of dying at age x having survived to age $x-1$ (the hazard rate) by q_x. Suppose that the insurance company forecasts mortality improvements by adjusting data from an actuarial table q_x^0 by multiplying by an exponential factor f^{x-x_0}, where x_0 is the current age of the annuitant and f is a scalar (which is less than unity if mortality improves over time and equal to unity in the case of no mortality improvement). This is one way by which the U.K. Institute of Actuaries' Continuous Mortality Investigation Bureau (CMIB) makes mortality adjustments.

In terms of the q_x, the unconditional probability of dying after $T > 0$ periods (conditional on having lived to age x_0) is

$$q_{T+x_0} \prod_{x=x_0}^{x_0+T-1} (1 - q_x),$$

where q_x is the conditional probability of dying at age x having survived to age $x-1$.

This unconditional probability is used in computing the expected value in Equation (2), so Equation (2) is equivalent to (if we also take into account the improvement factors)[2]

$$P = \sum_{T=1}^{\infty}\left[\frac{d}{r}\left[1 - (1 + r)^{-T}\right]\right] q_{T+x_0}^0 f^T \prod_{x=x_0}^{x_0+T-1} (1 - q_x^0 f^{x-x_0}). \tag{3}$$

[2] This expression takes the value of the bond if it pays out for T years and multiplies by the *unconditional* probability of surviving for T years and then sums over all possible values of T. It is equivalent to the following expression using *conditional* probabilities:

$$P = \sum_{T=1}^{\infty} d \quad q_{T+x_0}^0 f^T (1 + r)^{-T}.$$

Errors in the adjustment factor f can have a large impact on Equation (3). Historical evidence on mortality forecasts suggests that forecast errors of 15 to 20 percent in f for intervals of ten or more years ahead are not uncommon (see Table 18.7 in MacDonald [1996]). Another indicator of the difficulty in forecasting mortality improvements is that historical values of these improvement factors are not constants; they differ considerably between men and women and across ages and types of pensioners. For instance, the historical improvement rate for men aged 70 between 1967 and 1970 and between 1979 and 1982 was 0.74 for life office pensioners and 0.91 for immediate annuitants, i.e., those who purchase annuities voluntarily (MacDonald, 1996). The impact of such forecast errors on survival probabilities is significant. For example, assuming a 20-year improvement factor of 0.80 for a 65-year-old man and forecast errors for mortality rates of up to 10 percent over a ten-year period, then using the PMA80 mortality tables (constructed to reflect the mortality experiences of males who are obliged to buy pension annuities), the forecast probability of a 65-year-old man living to 85 ranges from 33.7 percent to 43.8 percent, while that of his living to 95 ranges from 5.3 percent to 15.3 percent.[3]

To determine the effects of these forecast errors on annuity yields, the authors use Equation (2) to solve for the annuity yield d/P:

$$\frac{d}{P} = E\left[\frac{\left[1-(1+r)^{-T}\right]}{r}\right]^{-1}. \tag{4}$$

To compute the actuarially fair yield (i.e., with zero cost loading), the authors substitute in survival probabilities determined from standard actuarial tables into Equation (4). For a male aged 65 and a discount rate of r = 7 percent and a 20-year improvement factor of 0.80, the PMA80 tables and Equation (4) lead to an actuarially fair annuity yield of 10.6 percent, but forecast errors suggest it lies between 10.3 percent and 10.9 percent. Thus, the percentage difference in yields is about 5 percent.[4] For a woman, the PFA80 tables lead to an actuarially fair yield of 9.5 percent in the absence of mortality risk, but forecast errors suggest it lies between 9.2 percent and 9.7 percent—again a percentage difference of about 5 percent.

The effects of mortality forecast errors are more serious for escalating annuities because payments in the future will be higher than with flat annuities. For escalating annuities, Equation (4) becomes:

[3] The quoted improvement factors are in units of 20 years and need to be converted to 1-year factors. This is done in the authors' case as follows. For a 20-year improvement factor of 0.8, the 1-year improvement factor is $0.8^{\frac{1}{20}}$. The lower bound on the improvement factor is the 1-year improvement factor times $0.90^{\frac{1}{10}}$, while the upper bound is the 1-year improvement factor times $1.10^{\frac{1}{10}}$. This converts a 10 percent forecast error either way over 10 years to the appropriate 1-year forecast error.

[4] That is $\frac{10.9-10.3}{10.6}$.

$$\frac{d}{P} = E\left[(1+\pi)\frac{\left[1+\left(\frac{1+r}{1+\pi}\right)^{-T}\right]}{r-\pi}\right]^{-1}, \tag{5}$$

where π is the uprating factor. For example, with an annuity escalating at 4 percent per annum, the percentage difference between upper and lower bounds[5] rises considerably: To 9.5 percent for women and to 8.9 percent for men.

Given the significance of mortality forecasts for insurance company profitability, it is not surprising that cost loadings to cover mortality risk are built into prices or that some insurers simply offer uncompetitive annuity rates, thereby effectively staying out of the market.[6] Insurance companies cannot at present reduce these loadings without taking on unreasonable risks; indeed, anecdotal evidence for the U.K. indicates that the failure of some insurance companies to accurately predict improvements in mortality has led to serious problems among suppliers of deferred annuities, which are even more susceptible to mortality risk than immediate annuities.[7] Similar calculations for a 20-year deferred annuity escalating at 4 percent per annum to be received by a woman when she reaches the age of 65 suggest a range of about 22 percent in annuity yields under different mortality forecasts. These business considerations suggest why the market for deferred annuities in the U.K. is relatively thin.[8] Deferred annuities are particularly important in the case in which a defined benefit plan is wound up, say, as a result of the insolvency of the sponsoring company and also, potentially, for early leavers.

Because the private sector is less able to absorb the aggregate risks associated with mortality forecast errors than is the government, the authors will examine what the government can do to help alleviate this problem and reduce costs to annuitants.[9]

---

5 The authors use the term "bound" to represent outcomes with typical historical forecast errors.

6 The difference between the best and worst annuity yields can be as much as 30 percent over all the companies offering annuities in the U.K., yet there is almost no variation in the type of policy offered. Even across the top ten providers (which account for the bulk of the market), the differences can be substantial and there is a distinct absence of competition among the remaining 230 providers. Because the purchase of retirement annuities is mandatory in the U.K., and because many thousands of workers retire every week, it is easy for companies wanting to attract new business to do so by temporarily offering the best annuity yield. However, there must invariably come a point when some of these companies either will have written their quota or will decide to concentrate on other types of business that are subject to smaller long-term risks.

7 Anecdotal evidence for the U.K. also suggests that currently life companies are having to make annuity payments for two years longer than originally anticipated.

8 Only £10 million in single premium deferred annuities were issued in 1996, about one-eightieth of the level of single premium immediate annuities sold in that year (Association of British Insurers, 1997 [Table 12]).

9 The government has a major responsibility here, because its public health campaigns are aimed directly at improving the mortality of the whole population; this has important implications for annuity provision.

Introducing Survivor Bonds

When insurance companies write annuities, they use the premiums collected to buy *matching assets*, that is, assets whose cash payments match as closely as possible the anticipated pattern of payouts on the liabilities that they face. In the case of level annuities, they invest principally in fixed-income bonds. In the case of index-linked annuities, they hold index-linked bonds; no insurance company would be prepared to write index-linked annuities if it could not lay off the resulting inflation risk through the purchase of an index-linked bond issued by, say, the government or a utility. However, insurance companies face one risk for which *no* matching assets exist: mortality risk.

A simple solution to the problem of mortality risk would be for the government to issue *survivor* (or *life annuity*) *bonds*, that is, bonds whose future coupon payments depend on the percentage of the whole population of retirement age (say 65) on the issue date still alive on the future coupon payment dates. For a bond issued in 2000, for example, the coupon in 2020 would be proportional to the fraction of 65-year-olds in the population who have lived to age 85. The coupon is therefore directly proportional to the amount an insurance company needs to pay out as an annuity to the average individual with an average pension. A new tranche of bonds would be issued on a unisex basis at the start of every year. The issue price would be determined by the Government Actuary.[10] The bonds would be traded on the open market and could be resold. Large occupational plans that also bear aggregate mortality risk would similarly be natural purchasers of such bonds.

The Implications of Survivor Bonds

The authors will examine the implications of introducing survivor bonds both for annuity providers and for the government.

Survivor bonds aim to lower the costs of retirement provision for the average pensioner. This is because they are designed to hedge *aggregate* mortality risk. They are not designed to hedge *select* mortality risks. There are some key select risks to take into account:[11] (1) pensioner annuitants are likely to live longer than the average of the population of the same age and (2) given that an insurance company is underwriting a finite sample of lives, the characteristics of any particular insurance company's pool of annuitants may differ from those of the pensioner annuitant population as a whole. For example, women and wealthy pensioner annuitants with large lump sums to annuitize tend to live longer than the average pensioner annuitant. The bonds the authors propose eliminate only the risk associated with aggregate mortality improvements for the average of the whole population and do not eliminate select risks such as those associ-

[10] The authors do not envision any major problems with determining the issue price. The Government Actuary's Department could publish its underlying assumptions concerning mortality. The authors believe that the risk of underestimating mortality improvements is a lesser risk than that of underestimating inflation, and there appears to be no problem with determining the issue price of retail (consumer) price index-linked bonds. Similarly, there are bound to be measurement errors in determining the exact number of the original cohorts who die every year, but the authors do not believe that the order of magnitude of these errors is greater than that involved in determining the "true" inflation rate.

[11] For the U.K., Finkelstein and Poterba (1999) have quantified these select risks.

ated with gender or wealth.[12] In short, by minimizing aggregate risk, insurance companies would have the proper incentives to develop wide mortality pools and do what they do best: provide insurance against idiosyncratic risks.

For instance, if the mortality of the rich improves more than that of the poor and the insurance company chooses an equally weighted pool of rich and poor, the payouts by the insurance company will decline less rapidly than the coupon payments from the survivor bonds. The rate at which this happens depends on the differences between terms $q^0_{T+x_0} f^T \prod_{x=x_0}^{x_0+T-1} (1 - q^0_x f^{x-x_0})$ in Equation (3) for the insurance company's own annuitant pool and those for the population as a whole. The qs and fs will be lower for the insurance company's annuitant pool and the forecast errors higher than for the population as a whole. The implication of this is that insurance companies that choose to have mortality pools different from the population at large continue to bear select mortality risk, but that is a commercial decision.

However, the important point to recognize is that no obvious matching assets exist for the select mortality risks that are assumed by annuity providers once they hold survivor bonds to hedge aggregate mortality risk. The provider will hedge these select risks by offering *lower* annuity rates to *all* annuitants. This disadvantages the average annuitants who are not members of the select groups. The *only* way of dealing with this problem is to reduce the select mortality risks to zero by having the government make pension plans and pension annuities mandatory for all members of society. But this leads directly to questions about the role of government in pension provision.

Why should the government (and ultimately taxpayers) issue survivor bonds and absorb the risks associated with mortality fluctuations? One possible justification can be found in the Arrow-Lind Theorem (1970) on social risk-bearing, which shows that, by dispersing an aggregate risk across the population (of taxpayers) as a whole, the associated risk premium can be reduced to zero. The government could therefore issue survivor bonds at a lower yield (namely, the risk-free rate) than any private corporation could. The private corporation will have many fewer shareholders than there are taxpayers, and some of the shareholders may hold large blocks of shares that constitute a significant proportion of their net worth. These shareholders will demand a risk premium, whereas the government can act as a risk-neutral player. Another justification lies in the government's own public health campaigns that are aimed directly at improving the mortality of the whole population, and this has important implications for annuity provision by the private sector.[13] Similarly, the reform of social security pensions (which are themselves nontradeable survivor bonds)

---

[12] It would, of course, be possible for the government to choose the population of pensioner annuitants as "the population" for which it issues survivor bonds or to design the bond to meet the more specific circumstances of the pensioners (such as gender- or wealth-adjusted mortality rates). But it is unlikely that any government would agree to do this, because this would involve a substantial cross-subsidy from the general taxpayer.

[13] The authors note that the introduction of survivor bonds perversely provides the government with an incentive to introduce measures that reduce longevity among the aged, e.g., reducing medical expenditure and tobacco taxes on the elderly.

and the transfer of pension provision from the public to the private sector would be greatly eased by the existence of tradeable survivor bonds.

By issuing survivor bonds, the government would be helping to complete markets. But is the government really the best organization to issue them? In practice, a key assumption underlying the Arrow-Lind Theorem may be violated. Real-world tax systems do not give each household an equal share of the tax burden, and some households may be affected directly by the variation in mortality risks. For example, if older households experience an improvement in life expectancy, their resources may be stretched as they try to finance a longer-than-expected retirement period. Similarly, the government would have to collect more taxes to pay the extra coupons on the survivor bonds, and when this happens, if the taxes are also levied on older workers, they face a tax burden that is correlated with other shocks to their well-being, thereby violating one of the conditions of the Arrow-Lind Theorem. In contrast, when an insurance company bears the risk, the shareholders have volunteered to bear the mortality risk. Some investors may be more risk-tolerant than the typical taxpayer, and this could provide a more efficient means of risk sharing. Nevertheless, while insurance company stockholders may indeed be more risk-tolerant than the average taxpayer, it remains the case that no insurance company in the world has yet issued a single tradeable survivor bond.

However, insurance companies would appear on the surface to be natural issuers of such bonds, because they are in a position to hedge mortality risk with their other products: greater longevity raises the payouts on life annuities but lowers them on endowment policies. However, in practice, endowment policies provide a poor hedge for life annuities, because mortality improvements are not spread evenly across ages, but rather are concentrated at greater ages. To illustrate, the percentage improvement in mortality between the PMA80 and PMA90 tables (based on mortality experience for United Kingdom male annuitants in 1980 and 1990 respectively) was 12 percent at age 35, 9 percent at age 55, 23 percent at age 75, and 20 percent at age 95. Perhaps this is the reason insurance companies have chosen not to issue survivor bonds, despite the fact that there are no restrictions anywhere in the world preventing them from doing so.

Some have suggested that, because the market will price the bonds with an implied hazard rate that will coincide with expected mortality improvements, the outcome should not differ from that of using a large insurance company that provides annuities even if it does not sell survivor bonds. However, the forecasts of the market may not coincide with those of a particular insurance company, and insurance companies can (and do) go bust, unlike governments (at least in advanced economies), so the outcome cannot be identical, except in the case of homogeneous expectations and zero bankruptcy costs.

Finally, the family itself provides an informal mechanism for the issuing of survivor bonds among different generations of the same family, as implied by Kotlikoff and Spivak (1981), but the breakdown of the family in many countries makes this an increasingly unreliable mechanism.

So we are left with the state as the only realistic issuer of survivor bonds.

CONCLUSION

In much the same way as governments in a number of countries helped pension funds insure against inflation by issuing index bonds, the issuance of survivor bonds would help mature pension funds insure against the uncertainties involving an increasingly gray population. The authors believe that the reduction in cost loadings on annuities could be substantial.[14] However, the authors also note that their proposal is not new: survivor bonds are virtually identical in structure to the 1759 Geneva Tontine Bond (see Cooper, 1972; and Jennings and Trout, 1982).

REFERENCES

Arrow, K., and R. Lind, 1970, Uncertainty and the Evaluation of Public Investment Decisions, *American Economic Review*, 60(3): 364-378.

Association of British Insurers, 1997, *Insurance Statistics Yearbook, 1986-1996* (London: Association of British Insurers).

Blake, D., 2000, *Financial Market Analysis* (Chichester, U.K.: Wiley).

Blake, D., and J. M. Orszag, 1997, Towards a Universal Funded Second Pension: A Submission for the 1997 Pensions Review Focusing on the Financial Aspects of the Provision of A Second Tier Funded Pension, Special report, The Pensions Institute, Birkbeck College, University of London.

Blake, D., and J. M. Orszag, 1998, Stakeholder Pensions: Exploiting Scale Economies, Special report, The Pensions Institute, Birkbeck College, University of London.

Cooper, R. W., 1972, *An Historical Analysis of the Tontine Principle*, Monograph 1, S. S. Huebner Foundation, The Wharton School, University of Pennsylvania, Philadelphia.

Finkelstein, A., and J. Poterba, 1999, Selection Effects in the Market for Individual Annuities: New Evidence from the United Kingdom, Economics Department Discussion paper, Massachusetts Institute of Technology, Cambridge.

Jennings, R. M., and A. P. Trout, 1982, *The Tontine: From the Reign of Louis XIV to the French Revolutionary Era*, Monograph 12, S. S. Huebner Foundation, The Wharton School, University of Pennsylvania, Philadelphia.

Kotlikoff, L. J., and A. Spivak, 1981, The Family as an Incomplete Annuities Market, *Journal of Political Economy*, 89(2): 372-391.

MacDonald, A., 1996, United Kingdom, in: A. MacDonald, ed., *The Second Actuarial Study of Mortality in Europe*, Groupe Consultatif des Associations D'Actuares des Pays des Communautes Europeennes.

[14] For example, U.S. studies (e.g., Mitchell et al., 1999; and Poterba and Warshawsky, 1998) found that the deduction from the actuarially fair value of an annuity for a 65-year-old U.S. male was 15 percent if the male was a typical member of the population as a whole (calculated using the mortality tables for the whole U.S. male population) and 3 percent if the male was typical of the population buying annuities voluntarily (calculated using the select mortality tables for male annuity purchasers), implying a 12 percent deduction for the greater mortality risk.

Mitchell, O. S., J. M. Poterba, M. J. Warshawsky, and J. R. Brown, 1999, New Evidence on the Money's Worth of Individual Annuities, *American Economic Review*, 89(5): 1299-1318.

Poterba, J. M., and M. J. Warshawsky, 1998, The Costs of Annuitizing Retirement Payouts, Working paper, Economics Department, Massachusetts Institute of Technology, Cambridge.