

ISSN 1749-8279



World Economy & Finance
Research Programme
Working Paper Series

WEF 0027

Annuities and Aggregate Mortality Uncertainty

Justin van de Ven
Martin Weale

National Institute of Economic and Social Research

March 2007

Annuities and Aggregate Mortality Uncertainty*

Justin van de Ven and Martin Weale
National Institute of Economic and Social Research,
2, Dean Trench Street,
London SW1P 3HE

10th March 2007

Abstract

This paper explores the effect of aggregate mortality risk on the pricing of annuities. It uses a two-period model; in the second period people face a constant but initially unknown risk of death. Old people can either carry the aggregate mortality risk for themselves or buy annuities which are sold by young people. A market-clearing price for such annuities is established. It is found that old people would, given the choice, decide to carry a considerable part of aggregate mortality risk for themselves.

*This work has been supported by the ESRC under its World Economy and Finance Programme. Please treat the results presented here as preliminary.

1 Introduction

Many of the issues associated with annuity markets have been well-researched in recent years, as reported in the comprehensive survey by Davidoff, Brown, and Diamond (2005). Much of this research has focussed upon the merits of conventional annuities, with the general finding that there is likely to be a substantial role for annuitisation even in context of heterogeneous preferences for investment, the time stream of consumption, and bequests. In contrast, very little work has been conducted to understand the workings of conventional annuities markets with regard to aggregate mortality risk. This is surprising, given the role played by mortality risk in the *a priori* justification for conventional annuities. Here, we suggest an appropriate economic model for pricing conventional annuities, and use the model to project plausible values for the annuity risk premium attributable to uncertainty regarding individual and cohort specific life expectancy.

To clarify terms, aggregate mortality risk refers to the uncertainty that is associated with the average longevity of a birth cohort. Individual mortality risk, by contrast, is the uncertainty of the precise time of death of an individual, relative to their respective birth cohort. Individuals from a given birth cohort could protect themselves against individual mortality risk through the use of a mortality-adjusted annuity, which pays an annual dividend that depends upon the realised mortality rate of the cohort¹. Purchasers of such a mortality-adjusted annuity will, nevertheless remain exposed to income uncertainty to the extent that cohort mortality is uncertain. This differs from a conventional life annuity, which pays a guaranteed income to purchasers until death.²

The topic of aggregate mortality uncertainty was not discussed in the survey by Davidoff, Brown, and Diamond (2005). While Bohn (2005) does air the question of uncertain longevity, he limits himself to scenario analysis – the effect of a permanent increase in longevity during retirement – and does not address the question of risk in its more conventional sense. The problem is that the true mortality rate of any particular cohort is known only *ex post*. When an insurance company sells an annuity it contracts to pay a regular income calculated on the basis of a forecast of the mortality rate. Like any forecast, projections of mortality rates are subject to error, and any insurance

¹Conceptually a mortality-adjusted annuity is not very different from a tontine. However historically tontines have operated by pooling the contributions of the participants with the last survivor receiving the whole amount and the other participants receiving nothing.

²We are not aware of any contemporary sellers of mortality-adjusted annuities in the United Kingdom but it is likely that there would be benefits to filling this gap in financial markets; certainly, if the premium associated with aggregate mortality risk is large, then one might expect that annuitants would value the opportunity to carry at least some of the aggregate mortality risk themselves. Annuities are available which invest the funds in the stock market and which therefore pay out on a with profits basis; typically with these the annuitants do carry aggregate mortality risk.

company selling conventional annuities therefore carries mortality risk. Casual discussion in the United Kingdom has suggested that this risk may be substantial. Banks and Blundell (2005) make the point that “unanticipated increases in longevity put pressure on all forms of pension systems” although they do not offer any view on the scale of this pressure³. Hardy (2005) states that “Even actuaries recognize that longer life is a good thing – but, to the extent that it is unanticipated, it is also an enormous problem for the managers of annuity portfolios”, with the implication that an annuitant should expect to pay a substantial risk premium on an annuity. On the other hand studies such as Finkelstein and Poterba (2002) argue that, after allowing for adverse selection, annuity pricing is close to actuarially fair before making any allowance for a risk premium. This last observation implies that either the risk premium cannot be very large or that, once it is taken into account, conventional annuities are substantially mis-priced.

One clue to the pricing of mortality risk might be offered by financial markets. In November 2004 the European Investment Bank (EIB) issued a mortality bond following an earlier issue in December 2003 by the insurance company Swiss Re. A mortality bond is a loan stock whose payout depends on the mortality rate of a specific cohort and in the case of the EIB issue, the mortality of the cohort of men in the United Kingdom aged 65 in 2003. The pay-out received by purchasers of the stock increases proportionately to the longevity of the defined cohort, thus providing a hedge against aggregate mortality risk. Friedberg and Webb (????) explore the pricing of the EIB mortality bond in the context of the consumption capital asset pricing model (CCAPM). Despite their observation that conventional annuities expose insurance companies to “substantial risk”, they find on the basis of their model that the cost of hedging aggregate mortality risk should be very low.

The CCAPM prices assets taking into account correlations between associated investment returns and consumption needs. This is an appealing approach for pricing conventional annuities as it reflects the insurance motive. The CCAPM does not, however, provide an explanation of pricing in terms of demand and supply. Furthermore, the assumption of an infinitely lived representative consumer by the CCAPM sits uncomfortably alongside the uncertain mortality that influences demand for conventional annuities in practice.⁴

Annuities represent a transaction, where one group of consumers (the elderly) divest themselves of their aggregate mortality risk by purchasing insurance from another group (the young). In this

³In the United Kingdom there is a separate problem that mortality rates have been systematically over-predicted by the Government Actuary. Here, however, we focus on shocks relative to unbiased forecasts.

⁴The model could be extended to address aggregate mortality risk using the approach suggested by Yaari (1965) in which all consumers have the same mortality rate independent of their age and where, therefore, the notion of a representative consumer can be retained, as Blanchard (1985) does in his analysis of fiscal policy with finite horizons.

paper we consequently explore the pricing of aggregate mortality risk in an overlapping-generations model. In the model young people have no risk of death but, beyond a threshold age when they become old, they face a constant risk of death. Old people buy conventional annuities from young people (or from insurance companies whose shares are owned by young people) at a price which balances the willingness of the young to carry aggregate mortality risk with the desire of the old to divest themselves of it. The old and the young have different attitudes to the mortality risk of the old because they are of different ages; young people can adjust their consumption while still young in response to the gains or losses that they experience from carrying the mortality risk of the old. The fact that different generations are affected differently by shocks to the mortality rate of any particular cohort makes transactions in mortality risk possible.

We begin by discussing the way in which mortality rates for men have changed in the last sixty years or so. We proceed to set out a modelling framework to assess the impact of aggregate mortality risk on the pricing of conventional annuities and we use this to assess the risk margins which we might expect to observe. Finally we draw conclusions.

2 Mortality Risk and the Evolution of Mortality Rates in the United Kingdom

A time-series econometrician might think that there are two separate issues present here. Systematic under-prediction might be indicative of bias, while forecast errors arise simply because the future is uncertain; with good forecasting methods, therefore, the bias present might therefore be eliminated.

Reality is less clear. The background to the forecasting error in the United Kingdom can be seen in figures 1 and 2. These graphs show the logarithmic differences of decade average death rates taken from the Human Mortality Database for the United Kingdom since 1930. The figures for the decade starting in 2000 are calculated from mortality data for 2000-2003 and divided by 0.7 to reflect the gap to the mid point of these to the mid-point of the 1990-99 decade is only 7 years. The rates for the four age bands follow similar patterns with a structural break occurring by the 1980s as a modest trend improvement became an accelerating trend. The sharpest improvements are with the younger two age groups where the mortality rates are rather low anyway. Thus the impact of these logarithmic changes on life expectancy at age sixty-five is not as substantial as a quick look at the graph might suggest. For very old men the acceleration in the reduction in mortality is less marked, as 2 shows.

These data have led some authors to conclude that there is a cohort effect present with people

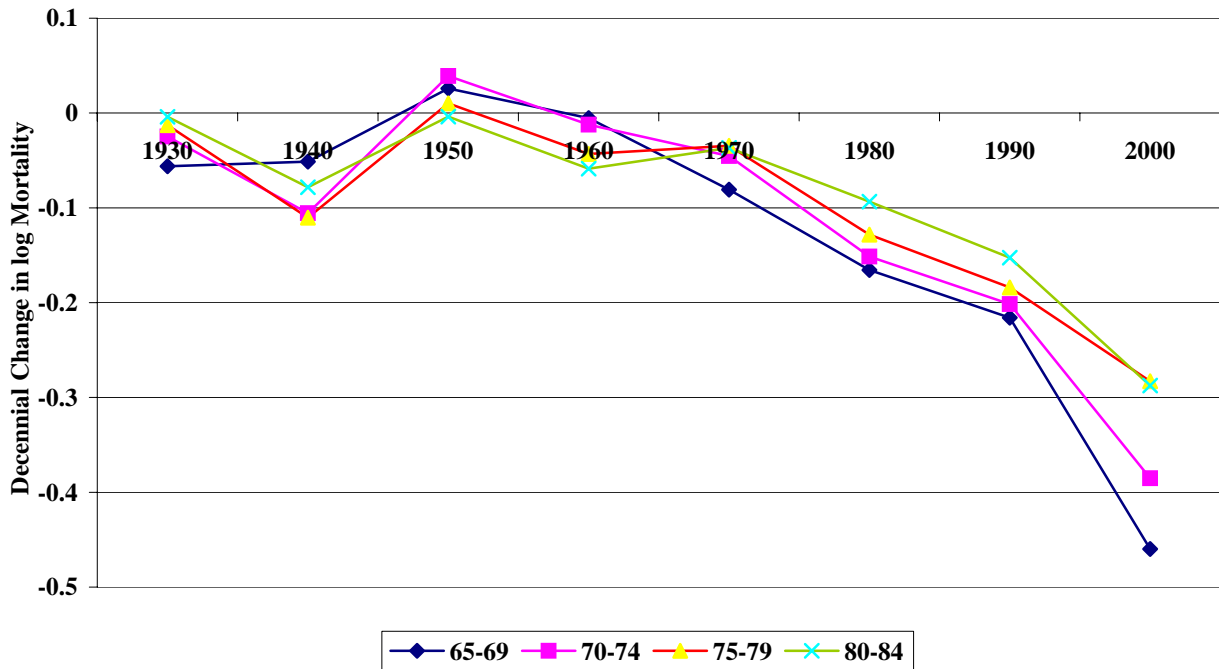


Figure 1: Decennial Changes in Log Mortality Rates: UK Men aged 65-84

currently in their seventies and early eighties experiencing much lower mortality rates than their predecessors. However, it remains to be seen whether that will continue to be the case as they reach extreme old age.

In forecasting mortality rates actuaries face the problem of deciding how long the rapid improvements shown in figure 1 will last- and the past pattern of mortality rates cannot be said to offer much of a guide to this; indeed it one looks further back- to 1840 the experience of the last thirty years or so stands out even more markedly. Not surprisingly views are divided, with some authors arguing that there is an upper limit to life expectancy⁵(Wilmoth and Horiuchi 1999)- with the implication not only that the rate of reduction in mortality rates must eventually stop, but that it must fall back towards zero, and other authors (Oeppen and Vaupel 2002) arguing that there is no sign of a limit being reached and plenty of evidence of assumed limits being broken. Of course the views need not be inconsistent; Oeppen and Vaubel are not arguing that life expectancy

⁵An upper limit to life expectancy does not imply an upper limit to human life. If the mortality rate for people aged 100+ settles at some figure below one, then an infinitely small number of people will live infinitely long. With a rate of 0.5, only one in a thousand centenarians would reach one hundred and ten, so it is not surprising that the statistical evidence is unclear.

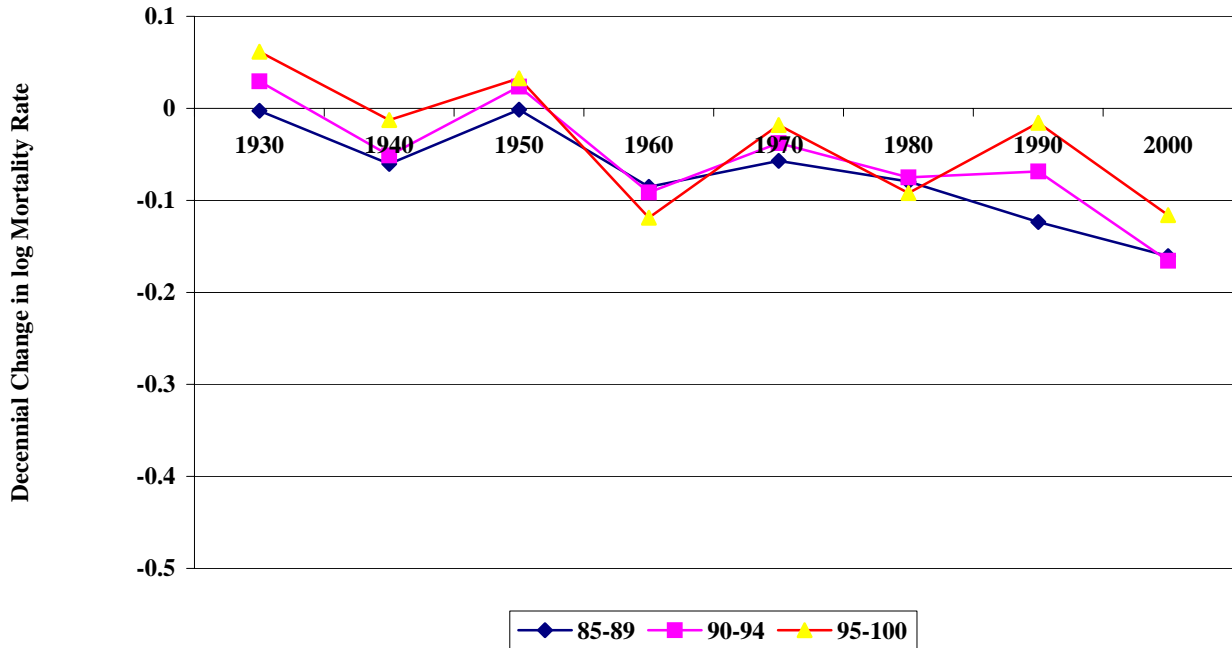


Figure 2: Decennial Changes in Log Mortality Rates: UK Men aged 85-104

will rise indefinitely but simply saying that, for practical purposes one should not assume that the improvement will stop.

The simplest method of projecting mortality rates is the Lee-Carter method (Lee and Carter 1992). This provides a structure in which a single factor, which can be modelled as a time series, is assumed to drive all mortality rates⁶. The uncertainty surrounding projections then depends on the uncertainty of the model parameters and the error term in the time series model for the mortality factor. Lee and Carter (1992) show that, for long-term projections of the type which concern us here, it is the uncertainty in the error term which dominates. Lee and Miller (2002) assess the performance of the model out of sample for the United States, looking at recursive forecasts for period (not cohort) life expectancy at birth for men and women combined produced

⁶The basic model underlying mortality is assumed to be

$$\log m(x, t) = a(x) + b(x)k(t) + \varepsilon(x, t)$$

where $m(x, t)$ is the mortality rate for someone of age x in year t . $a(x)$, $b(x)$ and $k(t)$ are model parameters and $\varepsilon(x, t)$ are error terms. The parameters are estimated to minimise $\sum_{x,t} \varepsilon^2(x, t)$ subject to the constraints that $\sum_t k(t) = 0$ and $\sum_x b(x) = 1$. Forecasts of mortality rates are produced by forecasts of k outside the sample.

from 1920 to 1998 and making the assumption that the mortality factor evolves as a random walk with drift. In the 1920s the median forecasts varied from 70 to 75 as compared to an outturn of 76.7 and an actual life expectancy between 55 and 60 in the 1920s. Despite this good performance, they acknowledge that the pattern of mortality decline has changed over the period, with declines in old age being more rapid than declines among young people.

Obviously, with models of this type it is possible also to produce estimates of uncertainty surrounding projections. However Pensions Commission (2005) considered a range of evidence and different projections for life expectancy, and suggested that, for men aged 65 in 2004 life expectancy lay in the range 17.6-19.9 years, while for a man aged 19 in 2004 it lay in the range 20-29 years. For our purpose, of using plausible estimates of mortality uncertainty to simulate a model we regard these ranges, which the Commission suggests could be seen as representing 90% confidence intervals, are perfectly adequate.

3 The Decision Framework

Our model structure focuses on two overlapping generations which we refer to as young and old. Young people begin their economic life at age 0 and all of them survive to age 1 when they become old. One period can be thought of as representing forty years. All old people in each cohort experience a cohort-specific mortality rate which is independent of age beyond the age of 1, but which is not known with certainty until after they have to make financial decisions for the period starting at age 1. The mortality rate in beyond age 1 for cohort i is defined as ρ_i . The analysis of conventional annuities presented by Yaari (1965) applies to the old generation with the qualification that at the age of 1 they can buy conventional annuities before their mortality rate is known. Alternatively they can buy mortality-adjusted annuities whose pay-out depends on their actual mortality rate once it is established. The interest rate in the model is r and the rate of discount of future utility is also assumed to be r .

Cohort i faces three decision points in our model.

1. At age 0 young people in cohort, i , receive an endowment of £ 1. They can sell conventional annuities to old people from cohort $i - 1$ who have now reached the age of 1. However, they will choose to do so only at a price which reflects the uncertainty surrounding the mortality of cohort $i - 1$; the expected mortality rate of cohort $i - 1$ is ρ_{i-1}^e . We denote by π_{i-1} the ratio of the actual annuity pay-out to that which would be made by a fair annuity without uncertainty. Thus an annuity bought by an old person in cohort $i - 1$ at age 1 pays a dividend

of $\pi_{i-1}(r + \rho_{i-1}^e)$. In the absence of uncertainty a fair annuity bought for £1 would yield a dividend of $r + \rho_{i-1}$. The capital value once the uncertainty is revealed is consequently $\pounds \frac{\pi_{i-1}(r + \rho_{i-1}^e)}{r + \rho_i}$ per £ 1 invested. The expected capital gain is therefore⁷

$$\pounds 1 - \pi_{i-1}(r + \rho_{i-1}^e)E\left(\frac{1}{r + \rho_{i-1}}\right).$$

Given the pay-out ratio in the market and the uncertainty surrounding the mortality rate of old people in cohort $i - 1$, young people in cohort i decide to sell conventional annuities to a value of $\pounds \phi'_i$.

2. At a time ε after this initial transaction, the actual mortality rate, ρ_{i-1} , of the old cohort becomes known. Given the realisation of ρ_{i-1} and thus the capital gains or losses that young people from cohort i have made, they choose their consumption rate c_i^y between the ages ε and 1. The assumption that the discount rate equals the interest rate means that this consumption rate is constant throughout the interval $[\varepsilon, 1]$.
3. At age 1 cohort i becomes old. It now holds wealth w_i^1 and has to decide what proportion of this, ϕ_i it wishes to annuitise. It makes this decision in the light of the uncertainty surrounding its own mortality rate, ρ_i , and in the face of the market pay-out ratio for conventional annuities, π_i . It buys a mortality-adjusted annuity with its remaining wealth⁸ giving a yield based on the as yet unknown mortality rate. This determines the income and consumption of the cohort when old, c_i^o .

We assume that ε is small so that we can ignore the consumption which takes place over the interval $[0, \varepsilon]$.

With an interest rate of r ,

$$w_i^1 = \left\{ 1 + \phi'_i \left(1 - \pi_{i-1} \frac{r + \rho_{i-1}^e}{r + \rho_{i-1}} \right) \right\} e^r - \frac{c_i^y}{r} (e^r - 1) \quad (1)$$

An old person with wealth w_i^1 at time 1 who invests a proportion ϕ_i in an annuity and the remainder in a mortality-adjusted annuity receives a level income of $y_i^o = \{(1 - \phi_i)(r + \rho_i) + \phi_i \pi_i (r + \rho_i^e)\} w_i^1$. Within the confines of our model, the results of Yaari (1965) and Davidoff, Brown, and Diamond (2005) imply that old people will invest all of their wealth in one or other of these products; they

⁷As the arithmetic mean is greater than the harmonic mean for any sequence of positive variables, it follows that for zero expected capital gain, $\pi_{i-1} < 1$. In other words, an annuity which is fair in the presence of uncertainty would seem unfair if assessed without taking that uncertainty into account.

⁸With our assumptions buying a tontine at time 1 is equivalent to buying an annuity at time $1 + \varepsilon$.

also imply that, provided the interest rate equals the discount rate, which we assume, these investors will choose consumption, $c_i^o = y_i^o$. Putting these expressions together allows us to write the budget constraint:

$$c_i^o = \{(1 - \phi_i)(r + \rho_i) + \pi_i \phi_i (r + \rho_i^e)\} w_i^o \quad (2)$$

$$= \{(1 - \phi_i)(r + \rho_i) + \pi_i \phi_i (r + \rho_i^e)\} \left\{ \left(1 + \phi_i' \left[1 - \pi_{i-1} \frac{r + \rho_{i-1}^e}{r + \rho_{i-1}} \right] \right) e^r - \frac{c_i^y}{r} (e^r - 1) \right\} \quad (3)$$

The last surviving cohort can sell conventional annuities to its predecessor but it cannot insure itself against its own longevity risk. For this cohort, cohort T , therefore, $\phi_T = 0$ and

$$c_T^o = (r + \rho_i) w_T^o \quad (4)$$

$$= (r + \rho_i) \left\{ \left(1 + \phi_T' \left[1 - \pi_{T-1} \frac{r + \rho_{T-1}^e}{r + \rho_{T-1}} \right] \right) e^r - \frac{c_T^y}{r} (e^r - 1) \right\} \quad (5)$$

Each consumer has an instantaneous utility function

$$u(c_i^z) = A + \frac{(c_i^z)^{1-\alpha}}{1-\alpha}; \quad z = y \text{ or } o$$

Here A is a variable which represents the joy of living and is large enough to ensure that $u(c_i^z) > 0$ for all plausible values of c_i^z . This condition has to be met for consumers to regard long life as a good thing rather than a burden but does not influence utility-maximising decisions..

4 Optimal Allocation

The solution to the model is found recursively. The presence of overlapping generations means that we have

Looking at the four decision points we work backwards, in common with the backward recursion used for solving dynamic programming problems. We define $V_{i,t}$ as the expected utility of someone in cohort i at time t on the assumption that they make optimal portfolio and consumption decisions and that they have a younger cohort with which to trade.

1. After the true mortality rate for old people is revealed at time $1 + \varepsilon$, there is no further uncertainty and

$$V_{i,1+\varepsilon} = \frac{u(c_i^o)}{r + \rho_i}$$

2. At time 1, unless $i = T$, the consumer has to choose how much wealth to annuitise in order to maximise expected utility. It is assumed that old people cannot be sellers of either mortality-adjusted annuities or conventional annuities; $0 \leq \phi_i \leq 1$ ($i < T$).

$$\begin{aligned}
V_{i,1}(\pi_i, w_i^1, \rho_{i-1}) &= \underset{0 \leq \phi_i \leq 1}{Max} E(V_{i,1+\varepsilon}) \\
&= \underset{0 \leq \phi_i \leq 1}{Max} \frac{(w_i^1)^{1-\alpha}}{1-\alpha} \\
&\quad \int_0^\infty \frac{\{(1-\phi_i)(r+\rho_i) + \pi_i \phi_i (r+\rho_i^e)\}^{1-\alpha}}{r+\rho_i} f(\rho_i, \rho_{i-1}) d\rho_i \\
&\quad + \int_0^\infty \frac{A}{r+\rho_i} f(\rho_i, \rho_{i-1}) d\rho_i
\end{aligned} \tag{6}$$

We represent the density function of ρ_i as depending on the values of ρ_{i-1} to take account of the persistence of mortality rates (See section 5). We stress that the solution is a function of the annuity pay-out, π_i .

If $i = T$ then there is no optimisation involved and

$$V_{T,1}(w_T^1, \rho_{T-1}) = \frac{(w_T^1)^{1-\alpha}}{1-\alpha} \int_0^\infty \frac{\{(r+\rho_T)\}^{1-\alpha}}{r+\rho_T} f(\rho_T, \rho_{T-1}) d\rho_T$$

3. At time ε the rate of consumption when young must be chosen. However no new information accrues in the interval $[\varepsilon, 1]$ and therefore

$$V_{i,\varepsilon}(\pi_i, w_i^\varepsilon, \rho_{i-1}) = \frac{A(1-e^{-r})}{r} + \underset{c_i^y}{Max} \left\{ \frac{(1-e^{-r})}{r(1-\alpha)} (c_i^y)^{1-\alpha} + e^{-r} V_{i,1}(\pi_i, w_i^1, \rho_{i-1}) \right\} \tag{7}$$

For the last cohort with $i = T$, $V_{T,1}(w_T^1, \rho_{T-1})$ replaces $V_{i,1}(\pi_i, w_i^1, \rho_{i-1})$

4. At time 0 the consumer must decide how far it is prepared to sell conventional annuities to the current elderly. π_i is unknown because it reflects the demand for conventional annuities. The homotheticity of our utility functions makes clear that this depends on the amount of wealth that consumers own when they reach age 1, w_i^1 . This in turn depends on the profits that they make from selling conventional annuities when aged 0 and thus on the realisation of ρ_{i-1} . ϕ_i' is found by solving

$$V_{i,0} = \underset{\phi_i'}{Max} E\{V_{i,\varepsilon}(\pi_i, w_i^\varepsilon, \rho_{i-1})\}$$

We set

$$w_i^\varepsilon = 1 + \phi_i' \left(1 - \pi_{i-1} \frac{r + \rho_{i-1}^e}{r + \rho_{i-1}} \right) \tag{8}$$

. The homothetic form of

$$V_{i,\varepsilon}(\pi_i, 1, \rho_{i-1}) = A \left\{ \frac{1 - e^{-r}}{r} + \int_0^\infty \frac{f(\rho_i, \rho_{i-1})}{r + \rho_i} d\rho_i \right\}$$

means that, with

$$U_i(\rho_{i-1}) = A \left\{ \frac{1 - e^{-r}}{r} + \int_0^\infty \frac{f(\rho_i, \rho_{i-1})}{r + \rho_i} d\rho_i \right\}$$

$$\begin{aligned} V_{i,0} &= \frac{Max}{\phi'_i} E \left\{ \left(1 + \phi'_i \left[1 - \pi_{i-1} \frac{r + \rho_{i-1}^e}{r + \rho_{i-1}} \right] \right)^{1-\alpha} \left(V_{i,\varepsilon}(\pi_i, 1, \rho_{i-1}) - U_i(\rho_{i-1}) \right) \right\} \\ &\quad + U_i \\ &= \frac{Max}{\phi'_i} \int_{\rho_{i-1}} \int_{\pi_i} \left(1 + \phi'_i \left[1 - \pi_{i-1} \frac{r + \rho_{i-1}^e}{r + \rho_{i-1}} \right] \right)^{1-\alpha} \\ &\quad (V_{i,\varepsilon}[\pi_i, 1, \rho_{i-1}] - U_i) h(\pi_i) g(\rho_{i-1}) d\pi_i d\rho_{i-1} \\ &\quad + U_i \end{aligned}$$

where $g(\rho_{i-1})$ is the density function of ρ_{i-1} conditional on the state of knowledge when cohort i is aged 0, i.e. on ρ_{i-2} , and $h(\pi_i)$ is the density function of π_i conditional on the same information. Since the realisation of ρ_{i-1} determines π_i through the market-clearing process we can write

$$\begin{aligned} V_{i,0} &= \frac{Max}{\phi'_i} \int_{\rho_{i-1}} \left(1 + \phi'_i \left[1 - \pi_{i-1} \frac{r + \rho_{i-1}^e}{r + \rho_{i-1}} \right] \right)^{1-\alpha} \\ &\quad \left\{ V_{i,\varepsilon}(\pi_i[\rho_{i-1}], 1, \rho_{i-1}) - U(\rho_{i-1}) \right\} f(\rho_{i-1}) d\rho_{i-1} \\ &\quad + \int_{\rho_{i-1}} U(\rho_{i-1}) d\rho_{i-1} \end{aligned} \tag{9}$$

5 Mortality Shocks

In common with empirical work on the topic, we assume that the mortality rate is log-normally distributed with density function $f(\rho_i)$. Thus

$$\log \rho_i = \frac{\gamma_1 + \nu_i}{\delta}$$

where ν_i is a normally distributed error term. Following the analysis of section 2 we assume that

$$\nu_i = \nu_{i-1} + v_i \quad \text{where } v_i \sim N(\gamma_2, 1), \quad Cov(v_i, v_j) = 0 \quad i \neq j$$

Working from ρ_{i-2} which is known with certainty at the time cohort i makes its first economic decisions, this gives

$$\begin{aligned}\log \rho_{i-1} &= \log \rho_{i-2} + \frac{v_{i-1}}{\delta} \\ \log \rho_i &= \log \rho_{i-2} + \frac{v_{i-1} + v_i}{\delta}\end{aligned}$$

Setting $\gamma_2 = -1/2\delta$ ensures that $E(\rho_i | \rho_{i-1}) = \rho_{i-1}$. The assumption of log-normality implies that we can use the Gaussian quadrature to compute expectations when we calculate the optimal choices.

In this paper we aim to produce a stylised overlapping generations model in order to identify the equilibrium price of aggregate mortality risk and to do this we make the following simple assumptions. We assume that life expectancy for men aged sixty-five is twenty years. For men aged below sixty-five the mortality rate is zero while for those aged sixty-five and over it does not vary with age; life expectancy at age sixty-five is therefore given by the reciprocal of the mortality rate once that is known. Log mortality follows a random walk process. We ignore the issue of drift although it would obviously be possible to simulate our model with any desired drift. We solve the model for a succession of periods and, noting that each period represents forty years, consider a value of γ_1 so that, in the initial period the expected mortality rate is set to two.

$$E \left\{ e^{\frac{\gamma_1 + v_i}{\delta}} \right\} = 2$$

We assume a standard deviation of 0.2, roughly equivalent to a standard deviation in life expectancy of two years so that

$$Var \left\{ e^{\frac{\gamma_1 + v_{i-1}}{\delta}} \right\} = 0.04$$

These conditions are met if $\gamma_1 = 6.9488$ and $\delta = 10.0249$ giving $\gamma_2 = 0.0499$. However, solution for the equilibrium pay-out ratio involves considering the effects of a range of different initial mortality rates, and in solving for the market-clearing pay-out we therefore use a range of values for γ_1 .

The standard deviation of the mortality rate of young men at age twenty-five is 0.285 giving a standard deviation of life expectancy of about 2.8 years. If we assume that the preferred Pension Commission figures represent 3.3 standard deviations (i.e. are taken from a normal distribution) then our assumption is a good fit for the mortality uncertainty faced by young men, but overstates that of old men. To the extent that this is the case, our simulations will overstate the aggregate mortality risk faced by old people and thus also their desire to protect themselves from it.

6 The Market Solution

We set out here demand and supply curves for conventional annuities.

6.1 The Demand for Annuities

The demand curve for conventional annuities is given as the value of ϕ_i which optimises equation (6). The homotheticity of the utility function implies that this is independent of w_i^o . Given our specification of ρ_i the integral can be evaluated using Gaussian quadrature; we use five abscissae with the assumption that $\nu_{i-1} = 0$. The optimum is then found using *fminsearch* in MATLAB. However, we note a tendency to over-annuitisation. With $\pi_i = 1$ we show in appendix A working to a second-order approximation, that $\phi_i = \frac{\alpha+1}{\alpha}$. Since this is independent of the magnitude of $Var(\rho_i)$ provided the latter is positive it is due to some fact other than the consequences of the differences between arithmetic and harmonic means noted above.

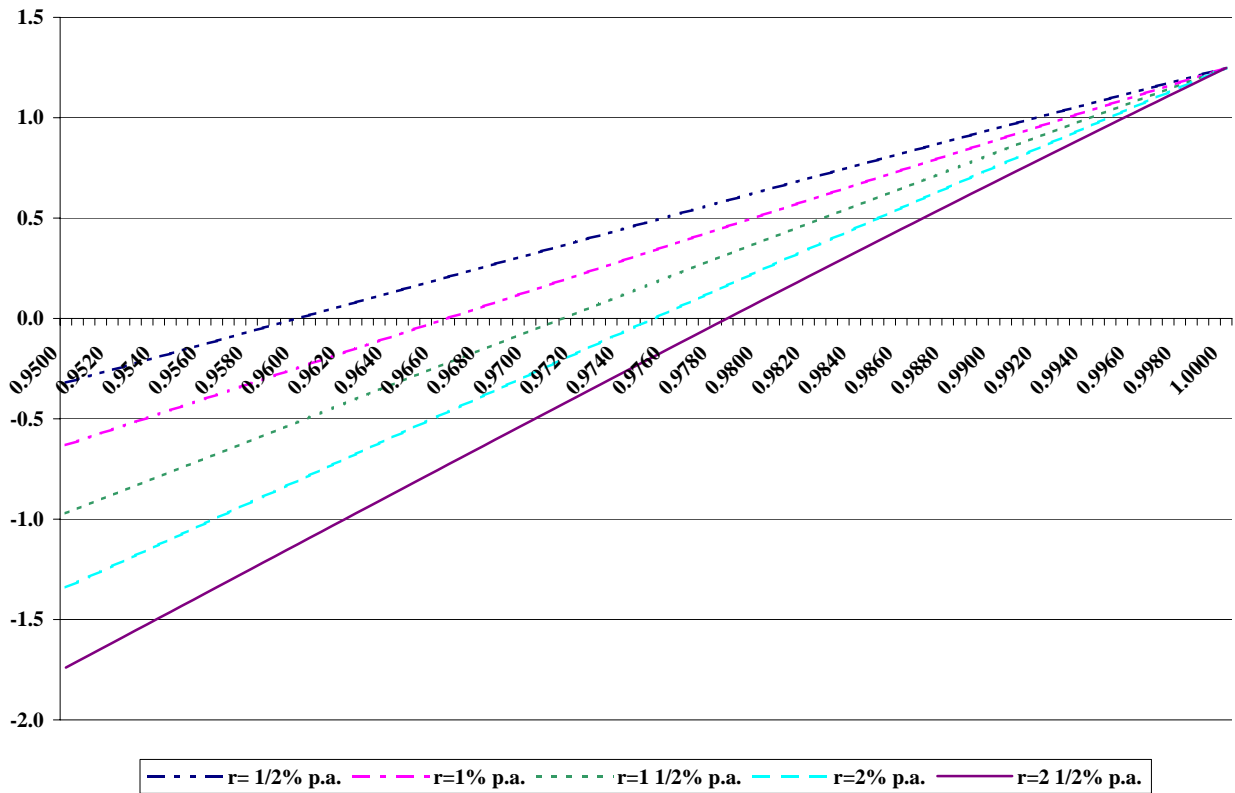


Figure 3: Demand for Annuities: Relative Risk Aversion = 4

The proportionate uncertainty in consumption arising from uncertainty in ρ_i is decreasing as a function of r . If the interest rate is low then the overall return on a mortality-adjusted annuity,

or on funds annuitised at age $1+\varepsilon$ is dominated by ρ_i ; it follows that shocks to ρ_i will have a large proportionate impact and people will be keen to buy conventional annuities even if the pay-out is poor. If r is large, then both ρ_i and shocks to ρ_i are likely to be less important. As a consequence, at any given value of $\pi_i < 1$ the proportion of w_i^1 which old people wish to annuitise will be a declining function of r . We illustrate demand curves for five values of r between $\frac{1}{2}\%$ p.a. and $2\frac{1}{2}\%$ p.a. in figure 1. The calculations are carried out on the assumption that an age of 1 represents forty years, so that r is forty times its annual rate. We use a value of $\alpha = 4$. This is towards the upper range of plausible estimates for this value; we choose a value which implies a high degree of risk aversion because the aim of this paper is, within the framework of a structural model, to look at the effects of risk aversion on the supply of and demand for conventional annuities.

The curves confirm that the steepness of the demand curve is a function of the interest rate, for the reasons we discussed above. At low annuity rates it is not surprising that old people would like to be net sellers rather than net buyers of conventional annuities. But the chart suggests that, for an annuity market to exist when investors can choose between conventional annuities and mortality-adjusted annuities, the value of π_i cannot be very far below 1. If sellers of conventional annuities need to make a substantial charge for risk, then there should not be an annuity market, and its existence can be explained only as a result of the legal requirement for pension funds to be used to purchase conventional annuities rather than mortality-adjusted annuities. If mortality-adjusted annuities are not available, the demand curve for conventional annuities is vertical.

6.2 The Supply of Conventional Annuities

While the demand curve for conventional annuities can be drawn unambiguously, the supply curve depends on the uncertainty that is associated with future pay-out rates. The homotheticity of the utility function implies that wealth w_i^1 when reaching age 1 is an increasing function of the profit realised on the sale of conventional annuities when young. Since the proportion of wealth annuitised at any pay-out ratio is independent of the amount of wealth, a high level of wealth will raise total demand for conventional annuities and thus depress the pay-out ratio. Thus the response of the pay-out ratio dampens, at least to some extent, the effect of uncertain returns from the sale of conventional annuities, compared to a situation where the pay-out ratio is unresponsive. To clarify the discussion here, we assume that the annuity pay-out consumers expect when old will be the same as that used to derive the relationship between current pay-out and supply; in other words $g(\pi_i) = \delta(\pi_i - \pi_{i-1})$ where $\delta(x)$ is the delta function with $\delta(x) = 0$ if $x \neq 0$ and $\int_{-\infty}^{\infty} \delta(x) = 1$, although when we subsequently solve for the market equilibrium, we take proper account of the role

of uncertainty in the subsequent pay-out ratio on the supply of annuities.

We find again that the steepness of the curve is increasing in the interest rate, showing in figure 4 curves constructed when mortality-adjusted annuities are assumed to be available. These curves once more reflect the fact that, with high interest rates mortality risk is quantitatively less important than with low interest rates. In consequence the proportion of wealth people are prepared to annuitise rises more steeply as the pay-out on the annuity falls. With a pay-out of 1 young people want to be buyers rather than sellers of conventional annuities (selling a negative proportion of initial wealth). This is a consequence of the fact that, with uncertainty, the fair pay-out ratio is below 1. Once again the higher is the interest rate, the less important is mortality uncertainty and thus the fair pay-out ratio is closer to 1.

When we look at the effects of persistence of mortality shocks we see the rather different picture shown in figure 4. When young investors lose money because their parents live for longer than expected, the persistence of mortality rates mean that, in addition they have a longer expected period of retirement to support. Thus the financial effect of losses arising on the sale of conventional annuities is magnified compared with the situation where mortality rates are independent. It is not very surprising that the supply curves are both shallower and shifted to the left as compared to the position where mortality rates are independent. The effect is at its most marked when the interest rate is very low for the reason identified earlier. When interest rates are very low the capitalised effects of shocks to mortality are greater than when they are higher.

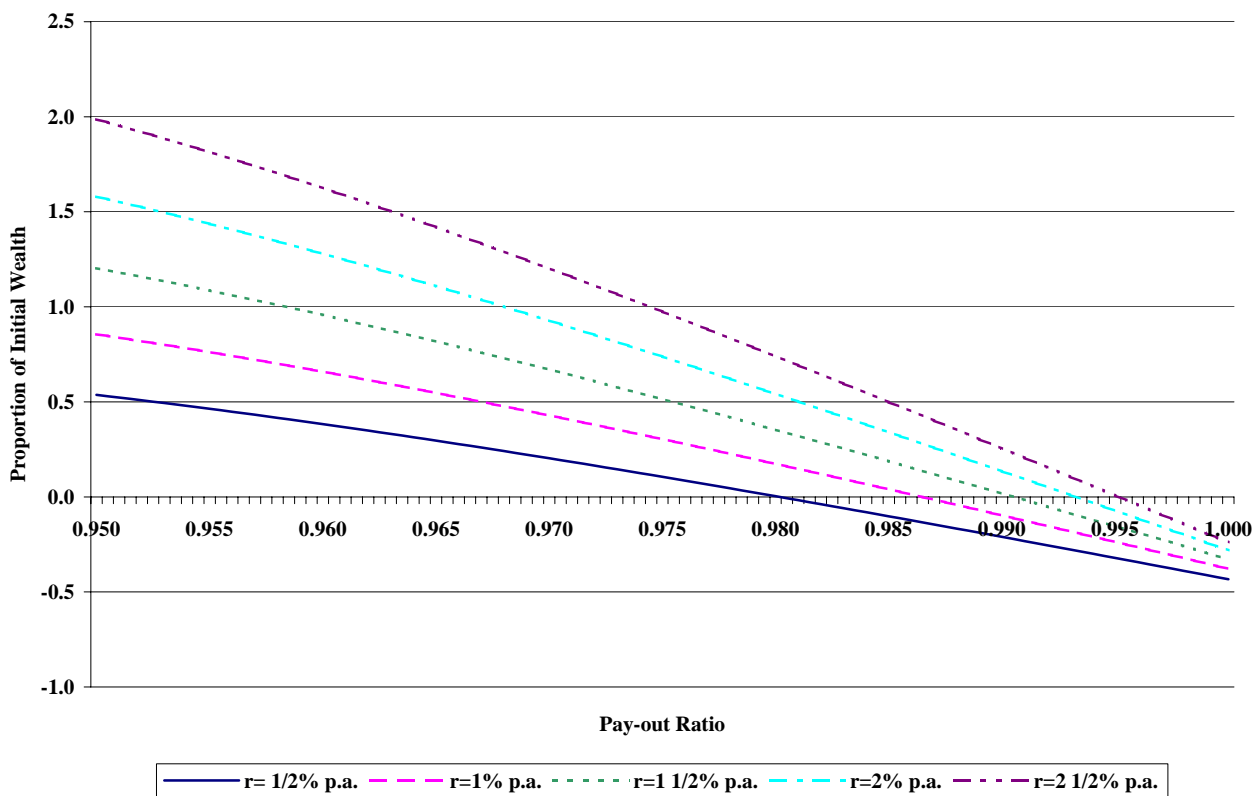


Figure 4: Supply of Annuities : Relative Risk Aversion = 4, Mortality Shocks follow Random Walk, Mortality-adjusted Annuities available

7 Market Equilibrium

7.1 Solution Method

As with any structural model of demand and supply, we can identify the market-clearing price- in this case we focus on its reciprocal, the pay-out ratio. We simulate a market-clearing process with the price adjusting so that the share of old people's wealth which young people are prepared to annuitise is equal to the proportion that old people would like annuitised. The reciprocal of the pay-out ratio can be thought of as the risk premium associated with demographic uncertainty.

The model solution is recursive. We work backwards from the market equilibrium established when young cohort T sells conventional annuities to old cohort $T - 1$. Both demand and supply depend on the mortality of cohort $T - 1$ and the subsequent mortality of cohort T . These are in turn stochastic functions of the known mortality rate of cohort $T - 2$. The market equilibrium also depends on the wealth of cohort $T - 1$, w_{T-1}^1 since, this together with the pay-out rate, π_{T-1} determines the demand for conventional annuities.

The equilibrium condition which determines π_{T-1} is

$$\phi'_T(\pi_{T-1}, \rho_{T-1}, \rho_T) = \phi_{T-1}(\pi_{T-1}, \rho_{T-1}) w_{T-1}^1$$

where the mortality rates ρ_{T-1} and ρ_T have to be forecast. The densities of both of these are determined by ρ_{T-2} which is known. For specified values of this and w_{T-1}^1 we can iterate to find the market-clearing value of π_{T-1} . We do this for a range of values⁹ and thus set up a grid showing the market-clearing mark-down as a function of these two variables, $\pi_{T-1} = \pi_{T-1}(\rho_{w-2}, w_{T-1}^1)$.

We repeat the process for earlier cohorts, noting that the supply of conventional annuities by young people is, in principle, a function of the uncertain markdown that they expect to have to face when old. Thus the supply function is

$$\phi'_i(\pi_{i-1}, \rho_{i-1}, \rho_i, \pi_i) \quad i < T;$$

with π_i depending on the as yet unknown w_{i-1}^1 and also what is known about the mortality risk after ρ_{i-1} has been realised.

Equation (8) indicates how w_{i-1}^1 depends on the out-turn for ρ_{i-1} relative to its expectation conditional on the portfolio choice of the cohort and the value of π_{i-2} . Thus we solve for the optimum value of ϕ'_i in equation (9) using the Gaussian quadrature with, at any given value of π_{i-1} , the quadrature points of the disturbances to (log) mortality allowing us to evaluate w_i^1 and π_i

⁹Using a convergence criterion of $\text{Abs}\{\phi'_\omega(\pi_{\omega-1}, \rho_{\omega-1}, \rho_\omega) - \phi_{\omega-1}(\pi_{\omega-1}, \rho_{\omega-1}) w_{\omega-1}^1\} < 10^{-4}$

| Log Mortality Random Walk | |
|---------------------------|--------|
| r=1% p.a. | |
| $E(\rho_{i-1})$ | 2 |
| π_{i-1} | 0.9846 |
| ϕ'_i | 0.2053 |
| ϕ_{i-1} | 0.5031 |
| c_i^y | 0.9029 |

Table 1: Market Equilibrium

by interpolating the grid. This procedure can be used to generate grids of the form $\pi_{i-1}(\rho_{i-2}, w_{i-1}^1)$ $i < T$ and thus we can find the trajectory of market clearing mark-downs for any given sequence of mortality shocks.

7.2 Results

As the description of the solution method makes clear, the supply of conventional annuities is in principle affected by the uncertainty that young people face about the circumstances in which they will be able to buy annuities in the future. For this reason the market-clearing pay-out ratio on conventional annuities alters as one approaches the terminal period. We solved the model for eight periods. We found that, working to four decimal places, while the pay-out ratio differs between the penultimate period and the final period, it does not change any more as one works further back. Thus the results we present can be regarded as steady state results.

We show our results in table 1 for the case where the real rate of interest is assumed to be 1% p.a. and $\alpha = 4$. The choice of a low real interest rate is deliberate since it creates conditions in which the effects of mortality risk are more likely to be pronounced than if the interest rate were higher. In the solution we present below cohort i sells conventional annuities to cohort $i-1$, and the table shows the market-clearing mark-down on the pay-out in the absence of uncertainty, π_{i-1} , the propensities to invest in mortality-adjusted annuities when young and when old (ϕ'_i and ϕ_{i-1}), and the expected propensity to consume out of wealth when young, c_i^y . All of these figures are shown on the assumption that cohort $i-2$ faced its expected aggregate mortality rate of 2 so that, for cohort $i-1$, $\log \rho_{i-1} = \frac{\gamma_1 + \nu_{i-1}}{\delta}$ where, as defined earlier, ν_{i-1} is normally distributed with mean $-1/2\delta$ and unit standard deviation.

These results point to a market equilibrium in which the mark-down on the annuity rate which would be observed in the absence of uncertainty is just over 1 1/2%. In equilibrium old people invest just over half their retirement wealth in conventional annuities.

8 Conclusions

This paper has set out a framework for exploring the issue of aggregate mortality risk by creating an overlapping generations structure in which young people carry the aggregate mortality risk of the old people. This creates a natural market and the market-clearing payout on annuities can therefore be explored. If old people have the choice of carrying some of aggregate mortality risk for themselves, the demand curve slopes fairly steeply. Even at a very low interest rate demand for conventional annuities falls to zero once the pay-out ratio drops to 96% of what a fair annuity would pay out in the absence of risk. Although the supply curve is less steep, it is not surprising that the market equilibrium is reached without a large charge for aggregate mortality risk. This suggests that aggregate mortality risk should not be a substantial influence on observed annuity rates.

References

- BANKS, J., AND R. BLUNDELL (2005): “Private Pension Arrangements and Retirement in Britain,” *Fiscal Studies*, 26, 35–54.
- BLANCHARD, O. (1985): “Debt, Deficits and Finite Horizons,” *Journal of Political Economy*, 93, 223–247.
- BOHN, H. (2005): “Who bears what Risk? An Intergenerational Perspective,” www.econ.ucsb.edu/~bohn.
- DAVIDOFF, T., J. BROWN, AND P. DIAMOND (2005): “Annuities and Individual Welfare,” *American Economic Review*, 95, 1573–1590.
- FINKELSTEIN, A., AND J. M. POTERBA (2002): “Selection effects in the United Kingdom individual annuities market,” *Economic Journal*, 112, 28–50.
- FRIEDBERG, L., AND A. WEBB (????): “Life is Cheap: using Mortality Bonds to Hedge Aggregate Mortality Risk,” NBER Working Paper No 11984.
- HARDY, M. (2005): “A Matter of Life and Death,” Financial Engineering News: [www.fenews.com/fen46/topics act analysis/topcis-act-analysis.htm](http://www.fenews.com/fen46/topics%20act%20analysis/topcis-act-analysis.htm).
- LEE, R., AND L. CARTER (1992): “Modelling and Forecasting US Mortality,” *Journal of the American Statistical Association*, 87, 659–671.

LEE, R., AND T. MILLER (2002): “Evaluating the Performance of the Lee-Carter Method for Forecasting Mortality,” *Demography*, 38, 537–549.

OEPPEL, J., AND J. VAUPEL (2002): “Broken Limits to Life Expectancy,” *Science*, 296, 1029–1031.

PENSIONS COMMISSION (2005): “Second Report: Appendix E. Uncertainties in Life Expectancy Projections,” www.pensionscommission.org.uk/publications/2005/annrep/appendices/appendix-e.pdf.

WILMOTH, J., AND S. HORIUCHI (1999): “Rectangularization Revisited: Variability at Age of Death within Human Populations,” *Demography*, 36, 475–494.

YAARI, M. (1965): “Uncertain Life-time, Life Insurance and the Theory of the Consumer,” *Review of Economic Studies*, 32, 137–150.

A Over-annuitisation

It is simple to demonstrate that, in the absence of a risk premium old people tend to over-annuitise.

We set $r + \rho_i = R_i$ and $r + \rho_i^e = R_i^e$. Then

$$\begin{aligned} c_i^o &= \{(1 - \phi_i) R + \pi_i \phi_i R^*\} w_i \\ \frac{\partial c_i^o}{\partial \rho_i} &= (1 - \phi_i) w_i \\ \frac{\partial c_i^o}{\partial \phi_i} &= \pi_i R^* - R \end{aligned}$$

and

$$E(u_i^o) = u(c_i^o) + \frac{\sigma^2}{2} \left\{ \frac{2u(c_i^o)}{R^3} - \frac{2u'(c_i^o)}{R^2} (1 - \phi_i) w_i + \frac{u''(c_i^o)}{R} (1 - \phi_i)^2 w_i^2 \right\}$$

Maximising this with respect to ϕ_i

$$\begin{aligned} \frac{\partial E(u_i^o)}{\partial \phi_i} &= \left(u(c_i^o) + \frac{\sigma^2}{2} \left\{ \frac{2u(c_i^o)}{R^3} - \frac{2u'(c_i^o)}{R^2} (1 - \phi_i) w_i + \frac{u''(c_i^o)}{R} (1 - \phi_i)^2 w_i^2 \right\} \right) (\pi_i R^* - R) \\ &\quad + \frac{\sigma^2 u'(c_i^o)}{R^2} w_i - \frac{\sigma^2 u''(c_i^o)}{R} w_i^2 (1 - \phi_i) \end{aligned}$$

If $\pi_i = 1$ and $R^* = R$ then the first order condition is

$$\sigma^2 u'(c_i^o) = \sigma^2 u''(c_i^o) R w_i (1 - \phi_i) \tag{10}$$

With the utility function given by (??) this suggests the solution

$$\phi_i = \frac{\alpha + 1}{\alpha} > 1 \text{ with } \alpha > 0$$

which is independent of the mortality risk. The result appears to suggest that overannutization takes place even in the absence of aggregate mortality risk. However if $\sigma^2 = 0$, then it is of course not permissible to divide both sides of (10) by σ^2 . It does, however indicate that with values of π_i close to but below 1 there is likely to be over-annuitisation.