Mortality Risk and Pricing of Annuities
National Institute Discussion Paper No 322

Justin van de Ven, and Martin Weale*

July 20, 2009

Abstract

Practitioners identify an average capital charge of 2.7% in pension buyouts as the cost of bearing uncertainty over future rates of mortality. This charge implies that it would cost around £60bn, or 4% of GDP, to insure the existing system of private retirement savings in the UK against the risks posed by mortality rate uncertainty. Using a plausible specification of mortality uncertainty, it is shown that only annuitants with relative risk aversion clearly higher than 20 would be prepared to pay this. Our results consequently provide an explanation for the thinness of the pensions buyout market.

Keywords: Pooled Annuity, Aggregate Mortality Risk, Risk Aversion

JEL: D14, D91, J11

1 Introduction

The pensions buy-out market in the United Kingdom (Lane, Clark and Peacock 2008) levies an average capital charge of 2.7% due to the uncertainty over mortality rate projections. The total value of life assurance and pension funds in the UK at the end of 2007 was £2,186bn, most of which represents savings for retirement. These figures suggest that – at current market prices – it would cost around £60bn, or just over 4% of GDP, for the system of private retirement savings to divest itself of the uncertainty over future mortality rates; and this cost excludes the additional exposure of the public pensions system. Despite the substantial size of these costs, the pricing of annuities in context of uncertain mortality rates (Ahlo & Spencer 1997) has only recently attracted attention. The purpose of this article is to consider the extent to which risk aversion can help to explain the significant market charges that are currently associated with mortality rate uncertainty.

---

*NIESR. `mweale@niesr.ac.uk`

Studies of the uncertainty attached to mortality rate projections (aggregate mortality risk) have not, hitherto, considered the implications for pricing of annuities. Lee (1998) and Renshaw et al. (1996) have produced stochastic models of mortality. These make it possible, by means of stochastic simulation, to produce estimates of density functions of life expectancy, and also to evaluate the density function of eventual profits or losses arising from any particular exogenous profile of annuity payments. A number of authors have used such models to explore the implications of uncertain mortality rates for the insurance companies selling annuities. Khalaf-Allah et al. (2006) look at the distribution of the costs faced by an insurance company selling annuities. Olivieri & Pitacco (2008*a*) investigate the implications of aggregate mortality risk for the cost of capital faced by an insurance company, while Olivieri & Pitacco (2008*b*) examine the interaction between stochastic mortality and the capital required to meet solvency requirements.

We extend this literature by considering what price people would, given the choice, be prepared to pay for protection from the effects of aggregate mortality risk. The approach that we take is to contrast preferences over a pooled annuity fund, with preferences over a standard level life annuity. Piggott et al. (2005) discuss the operation of a pooled annuity fund, where the assets of decedents in the fund accrue to survivors. Unlike a level life annuity, which protects annuitants from both individual specific and aggregate mortality risk, a pooled annuity fund only provides protection from individual specific mortality risk. The money value of the welfare disparity between level life annuities and a pooled annuity fund – equal to the compensating variation of replacing a life annuity with a pooled annuity – consequently provides an upper bound to the amount that people would be willing to pay to insure themselves against the uncertainty associated with mortality rate projections.

Our study is closely related to that of Piggott *et al.*, who explore how annuity payments may be varied in the light of deterministic, but unanticipated, mortality shocks. In a similar vein, Stamos (2008) explored the impact of the initial number of participants in a pooled annuity fund on the risk to which annuitants were exposed, and found that even rather small pools with one hundred initial members could produce 90% of the asymptotic gains associated with a pool of infinite size. In contrast to these studies, however, we focus upon a pooled annuity in which the annual pay-out is chosen at the beginning of each year to maximise the expected lifetime utility of annuitants, subject to an intertemporal budget constraint on the pool. This is the arrangement that would be made if each cohort were to carry its own aggregate mortality risk, in contrast to the conventional annuities market where the risk is carried by the seller of an annuity. Allowing for uncertainty in mortality rate projections requires the use of dynamic programming methods to solve for the intertemporal pay-out rate of a pooled annuity. In this context, we show that annuitants who are concerned about the risks to future annuity payments will want the pooled fund to pay out, in its early stages, less than it would if they were indifferent to those risks. The low initial payment is a form of precautionary saving, and the annuity payment is able to rise over time as a stock of precautionary savings builds up. An annuitant who is tolerant of risk will, by contrast, choose a path with a high initial payment and therefore greater risk that subsequent annuity payments will be low if survival rates are higher than expected.

We compare the expected welfare available to a purchaser of a pooled annuity, with the welfare available to a purchaser of a conventional annuity. We examine how far the return on the latter can be reduced below what is actuarially fair before the pooled annuity becomes relatively more attractive (in expected welfare terms). This allows us to identify the maximum amount that annuitants are willing to pay for protection from aggregate mortality risk, and thus sets an upper limit to the charge that sellers of conventional annuities would be able to levy, if annuitants had the alternative option to purchase a pooled annuity. In turn, this makes it possible to form a view about the market feasibility of the associated charges that have been identified in the pensions buy-out market.

Section 2 describes the analytical framework that we use to explore the pricing of annuities in context of aggregate mortality risk. In section 3 we report the pay-out profiles on pooled annuities that are implied by our analytical framework and associated welfare effects, focussing in particular on the implied “willingness to pay” for insurance against aggregate mortality risk. A summary is provided in the conclusion.

2 Pooled Annuities When Mortality Rates are Uncertain

To focus discussion, we consider pooled annuities that do not offer partner protection, purchased by 65 year old men in the UK in 2006. The annuities with which we are concerned pay out a variable income at annual intervals, and we are interested in determining how the annuity pay-out rate influences welfare in the context of mortality rate uncertainty, and the associated implications for structuring the annuity stream.

The section begins by describing the specification that is assumed for the intertemporal evolution of mortality rates. We then describe the behavioural framework that is used to characterise preferences for annuities, and conclude with a detailed discussion of the practical steps that were followed to obtain the analytical results that are reported in Section 3.

2.1 A model of uncertain mortality rates

We adopt a model of mortality disturbances that focuses on the uncertainty surrounding the trend rate of decline in log mortality. Our analytical problem is simplified, relative to the literature cited in the introduction, by the fact that we are concerned with an individual cohort, as opposed to the financing problem that is faced by an insurance company which has sold annuities to people from different birth cohorts. Thus we do not need to make the distinction between cohort effects and time effects. We assume that the mortality rates of the cohort of interest, ρ_t at age t, follow the process described by

$$\begin{aligned} \log \rho_t &= \begin{cases} \min\{\log \rho_t^* + u_t\} & \text{if } \rho_{t-1} < 1.0, t < T \\ 0.0 & \text{otherwise} \end{cases} && (1) \\ u_t &= u_{t-1} + \theta_t + \varepsilon_t + v_{t-1} && (2) \\ v_t &= v_{t-1} + \eta_t && (3) \end{aligned}$$

where T is the maximum length of life, and $\varepsilon_t \sim N(-\frac{\sigma_\varepsilon^2}{2}, \sigma_\varepsilon^2)$ and $\eta_t \sim N(-\frac{\sigma_\eta^2}{2}, \sigma_\eta^2)$ are independently distributed, both through time and with respect to one another. The respective distributions assumed for the disturbance terms ensure that $E(e^{\varepsilon_t}) = E(e^{\eta_t}) = 1$.

In the discussion that follows the expectation operator $E(.)$ and the variance operator $Var(.)$ are evaluated on the information available at the start of age 65, which is the age at which the annuity contracts considered for analysis are struck. In choosing parameters for the model, there are two distinct effects that need to be addressed. First, we want to ensure that $E(e^{u_t}) = 1$. And second, with $E(e^{u_t}) = 1$, we want to ensure that $E(\rho_t) = \tilde{\rho}_t$, where we envisage $\tilde{\rho}_t$ to be a series of exogenously defined cohort mortality rates as provided, for example, by the UK Government Actuary. ρ_t^* is the sequence of time varying non-stochastic terms which ensures that $E(\rho_t) = \tilde{\rho}_t$ for all t.

We are interested in the intertemporal evolution of mortality rates from age 65. The specification of η_t implies that $Var(v_t) = (t-64)\sigma_\eta^2$, and $E(v_t) = -(t-64)\sigma_\eta^2/2$, so that

$E(e^{v_t}) = 1$. Given this, if we are to ensure that $E(e^{u_t}) = 1$, then θ_t must be computed as follows. From equation (2) we have[1]

$$\begin{aligned} E(u_t) &= E(u_{t-1}) + \theta_t - \frac{\sigma_\varepsilon^2}{2} - (t-65)\frac{\sigma_\eta^2}{2} && (4) \\ Var\left(u_t\right) &= (t-64)\,\sigma_\varepsilon^2 + \sum_{i=65}^{t} (t-i)^2\,\sigma_\eta^2 && (5) \end{aligned}$$

For, $E(e^{u_t}) = 1$ we require $E(u_t) = -Var(u_t)/2$. It consequently follows that

$$\theta_t = \frac{\sigma_\eta^2\left[(t-65)-(t-65)^2\right]}{2} \tag{6}$$

If the distribution of ρ_t were not truncated, then θ_t as defined by equation (6) would ensure that $E\left(\rho_t\right) = \rho_t^*$, and so we would set $\rho_t^* = \tilde{\rho}_t$. Given the truncation that is assumed for ρ, however, it is necessary to solve numerically for the sequence of ρ_t^* such that $E(\rho_t) = \tilde{\rho}_t$. We explain how this is done in the section that follows.

2.1.1 Selecting model parameters

We first set the standard errors of the disturbances to $\sigma_\varepsilon = 0.01$ and $\sigma_\eta = 0.03$ since this gives a plausible degree of overall uncertainty for our simulations reported below. The sequence θ_t is then given by equation (6). We then need to find the sequence ρ_t^* to deliver the UK's official cohort life expectancy projections for men aged 65 in 2006.[2] We set the initial guess for ρ_t^*, $\rho_t^{*0} = \tilde{\rho}_t$ (the official mortality rate estimates), for all t. This would ensure that $E\left(\rho_t\right) = \tilde{\rho}_t$, if the distribution in equation (1) were not truncated by a maximum mortality rate of 1. To correct for the truncated distribution we adopt the following procedure to amend ρ_t^*.

1. Given the model parameters defined above, including ρ_t^{*0}, we use Monte Carlo methods to simulate stochastic outcomes from the model of mortality rates given by equations (1) to (3).

2. The (arithmetic) mean values of the age specific mortality rates resulting from one hundred thousand simulations, $\rho_t^{0,1}$, were calculated.

---

[1] The variance term can be derived directly from the alternative specification; $u_t = \sum_{i=65}^{t} \theta_i + \varepsilon_i + (t-i)\,\eta_i$.

[2] We thank the Government Actuary for providing the data upon which our model of stochastic mortality rates is based.

3. We then calculated $\rho_t^{*1} = \rho_t^{*0}.\tilde{\rho}_t/\rho_t^{0,1}$.

Steps 1 to 3 were then repeated, with $\rho_t^{*k} = \rho_t^{*k-1}.\tilde{\rho}_t/\rho_t^{k-1,k}$, until the difference between ρ_t^{*k-1} and ρ_t^{*k} was sufficiently small (less than 10^{-6} after five iterations). The effect of this change is to increase marginally ρ_t^* over $\tilde{\rho}_t$. However, even at age 109 the impact is only just over 0.5 percentage points.

The values of ρ_t^*, σ_ε and σ_η that were assumed for the analysis give a mean life expectancy[3] at age 65 of 20.6 years (with a standard deviation of 1.16 years), as compared to the official point estimate of 20.4 years. The small difference between the two arises because life expectancy is not a linear function of the mortality shocks as we have specified them; we have ensured that our stochastic model matches expected mortality at each age, as opposed to the measure of overall life expectancy. We have a 90% confidence interval for life expectancy of 18.7 to 22.5 years. This compares to a range of 17.6-19.9 years estimated by the Pensions Commission (2005). Thus our model delivers an interval slightly wider than that identified by the Pensions Commission. However the fact that the range is higher, and indeed that the Pensions Commission range does not include the more recent official point estimates, demonstrates just how rapidly views on mortality have been revised in recent years. Figure 1 shows the dispersion of the proportion of the population of sixty-five year old men surviving to each age shown on the graph. The median line, which divides the 25-50% range from the 50-75% range is scarcely distinguishable from the mean.

2.2 The welfare effects of a pooled annuity

The problem we address has its roots in the analysis of consumer behaviour in the context of uncertainty. In the absence of uncertainty, economic theory suggests that the time profile of consumption which maximises intertemporal utility will depend on the degree of impatience – commonly represented by a discount rate – relative to the real return on capital. If the discount rate is below the real return on capital, then consumption should be expected to rise over time, and *vice versa.* Leland (1968) shows that a risk averse consumer who faces an uncertain return on savings will, for any expected return, tend to save more than would be the case in the absence of uncertainty, provided that their tolerance of risk is inversely proportion to their accrued wealth. Preferences of this type are referred to as exhibiting “prudence”, and the additional wealth that is accrued in

---

[3] Based on one hundred thousand simulations.

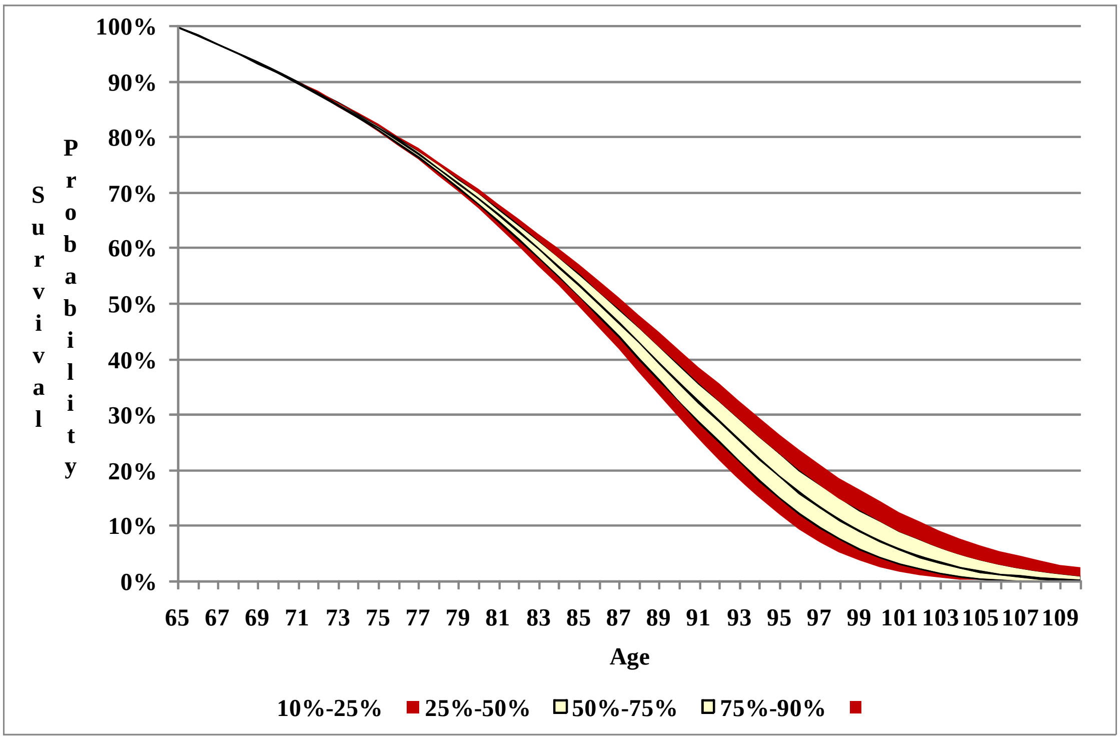

Figure 1: The Distribution of Survival Rates by Age from an Initial Population of Sixty-five Year Old Men in 2006

the context of uncertainty is referred to as precautionary savings.[4] The utility function that we consider here,

$$U_t = \sum_{j=t}^{T} \phi_{j-t,t} \delta^{j-t} u(c_t) \tag{7}$$

$$u(c_t) = \frac{c_t^{1-\alpha}}{1-\alpha} \quad , \alpha > 0 \tag{8}$$

$$\phi_{j-t,t} = \begin{cases} \prod_{k=t+1}^{j} (1-\rho_k) & \forall j > t \\ 1 & j = t \end{cases} \tag{9}$$

$$\rho_T = 1 \tag{10}$$

is just such a preference relation. Here, $\phi_{k,t}$ is the probability of surviving k years given survival to age t, δ is the discount factor equal to one minus the discount rate, c_t is consumption at age t, ρ_t is the period specific probability of mortality, and T is the maximum duration of life. This preference relation assumes constant relative risk aversion, so that, as people's wealth increases, their tolerance of any absolute degree of

---

[4] Preferences are described as exhibiting prudence if the sign on the third derivative on utility is positive ($u''' > 0$), which implies decreasing risk aversion ($-u''/u'$) with consumption.

uncertainty about future income increases.[5]

The use of a preference relation that is capable of reflecting precautionary saving is highly relevant to our problem. This is because a pooled annuity has the effect of sharing the assets of those who die in each year among the survivors, thereby raising the return on survivors' capital. Hence, when future rates of mortality are uncertain, so too are the implied future rates of return to wealth held in the form of a pooled annuity. The preferences that we assume consequently imply that, in the context of uncertainty over future mortality rates, consumers will prefer a pay-out path from a pooled annuity fund that increases with time, relative to the path that would be preferred in the absence of such uncertainty.

There is, however, an important qualification to this. We have assumed a maximum life-span of 110. Since the mortality rate is stochastic, there is plainly a risk of the entire cohort dying out before the age of 110. If this happens there will be no remaining survivors who might share out the remaining assets of the pooled annuity fund. From the perspective of the cohort, these funds are therefore lost. Once the risk of the cohort mortality rate rising to one becomes significant, in very extreme old age, the management of the annuity fund has to balance the risk of the fund's assets being wasted against the alternative risk of the cohort having lower than expected mortality. The preferred payout schedule will reflect the balance between these two concerns.

We now address the technical details of our optimisation problem. We assume that all annuitants have identical preferences, as represented by equation (7). We define the pay-out in any year from a pooled annuity as the amount that each annuitant would rationally choose to consume in the respective year, given their age and the aggregate fund available. This structure is important because it implies that the full benefits of annuitisation are delivered, subject to the annuitant carrying the aggregate mortality risk. This would not be the case if precautionary savings in response to uncertainty about future mortality rates were accrued outside of the pooled fund.

Let w_t define the value of the pooled annuity fund at the start of the year in which the annuitised cohort is aged t. Our optimisation problem can then be represented by the Bellman equation

$$V_t\left(w_t\right) = \max_{c_t}\left\{u\left(c_t\right) + \delta E_t\left[\left(1-\rho_t\right)V_{t+1}\left(w_{t+1}\right)\right]\right\} \tag{11}$$

---

[5] Pratt (1964) offers an account of the concepts of absolute and relative risk aversion.

where

$$w_{t+1} = \left(\frac{1+r}{1-\rho_t}\right)(w_t - c_t) \tag{12}$$

since it is assumed that the remaining assets of the members of the pooled annuity who die in year t accrue to the survivors. The uncertainty in ρ_t means that, for known values of the other variables, w_{t+1} is uncertain too.

There is a particular problem which arises if the probability that $\rho_t = 1$ is non-zero, i.e. that there is a finite risk that the cohort dies out completely in period t. Expression (12) implies that, unless $c_t = w_t$, w_{t+1} is unbounded as $\rho_t \to 1$. We define

$$\underset{\rho_t \to 1}{Lim} \; (1-\rho_t)\, V\left(\left(\frac{1+r}{1-\rho_t}\right)(w_t - c_t)\right) = 0$$

consistent with the obvious point that a cohort which dies out completely can derive no further welfare from any residual wealth and use this in the evaluation of the expectation in equation (11).

$V_t(w_t)$ is the total remaining expected life-time welfare as a function of w_t, on the assumption that optimal consumption choices are made, and conditional on being alive at the start of year t.[6] In the final period T, which represents the maximum possible life-span,

$$V_T(w_T) = \max_{c_T} u(c_T) \qquad c_T \leq w_T \tag{13}$$

with the trivial solution that all wealth is consumed, $c_T = w_T$.

Consider the problem in the penultimate period, $T-1$

$$\begin{aligned} V_{T-1}(w_{T-1}) &= \max_{c_{T-1}} \left\{ u(c_{T-1}) + \delta E_{T-1}\left[(1-\rho_{T-1})\, V_T\left(\frac{1+r}{1-\rho_{T-1}}[w_{T-1} - c_{T-1}]\right)\right]\right\} && (14) \\ &= \max_{c_{T-1}} \left\{ u(c_{T-1}) + \delta E_{T-1}\left[(1-\rho_{T-1})\, u\left(\frac{1+r}{1-\rho_{T-1}}[w_{T-1} - c_{T-1}]\right)\right]\right\} && (15) \end{aligned}$$

Defining $c_{T-1} = d_{T-1}w_{T-1}$, and substituting in for $u(c)$, we have,

$$V_{T-1}(w_{T-1}) = w_{T-1}^{1-\alpha}.\max_{d_{T-1}} \left\{ \frac{d_{T-1}^{1-\alpha}}{1-\alpha} + \frac{\delta}{1-\alpha} E_{T-1}\left[(1-\rho_{T-1})\left(\frac{1+r}{1-\rho_{T-1}}\right)^{1-\alpha}(1-d_{T-1})^{1-\alpha}\right]\right\} \tag{16}$$

---

[6] This is commonly referred to as the value function, which is a functional of the optimised consumption stream.

And generalising over all $65 \leq t \leq T-1$, given our focus on a pooled annuity purchased at age 65,

$$V_t(w_t) = w_t^{1-\alpha}.\max_{d_t}\left\{\frac{d_t^{1-\alpha}}{1-\alpha} + \frac{\delta}{1-\alpha}E_t\left[(1-\rho_t)V_{t+1}\left((1-d_t)\left(\frac{1+r}{1-\rho_t}\right)\right)\right]\right\}. \tag{17}$$

The solution to the decision making problem at any age t is given by the Euler condition,

$$d_t^{-\alpha} = \delta(1+r)(1-d_t)^{-\alpha}E_t\left[d_{t+1}^{-\alpha}\left(\frac{1+r}{1-\rho_t}\right)^{-\alpha}\right] \tag{18}$$

We set $\delta = 1/(1+r)$ so that – in the absence of uncertainty – the desired consumption stream is level over time. This gives

$$d_t^{-\alpha} = (1-d_t)^{-\alpha}E_t\left[d_{t+1}^{-\alpha}\left(\frac{1+r}{1-\rho_t}\right)^{-\alpha}\right] \tag{19}$$

so that

$$d_t = \frac{1}{1+\left(E_t\left[d_{t+1}^{-\alpha}\left(\frac{1+r}{1-\rho_t}\right)^{-\alpha}\right]\right)^{\frac{1}{\alpha}}} \tag{20}$$

This recursive equation provides the basis for evaluating the optimal annuity pay-out rate in each period.

2.3 Analytical approach

Despite the stylised nature of the stochastic process adopted for mortality rates as described in subsection 2.1, no analytical solution exists to the utility maximising problem outlined in subsection 2.2. Numerical methods based upon backward induction are consequently employed.

At each age, $65 \leq t \leq T$, a two dimensional grid is defined over the state space described by the potential values of u_{t-1} and v_{t-1}. In period T, the corner solution $d_T = 1$ is taken as given. We then work backward from the start of period $T-1$, solving for the decision d_{T-1} that is described by equation (20) at each intersection of the $T-1$ grid in (u_{T-2}, v_{T-2}), which is assumed to be known at the start of period $T-1$. The evaluation of equation (20) is complicated by the expectation in the denominator of equation (20):

$$\begin{aligned} E_{T-1}\left[d_T^{-\alpha}\left(\frac{1+r}{1-\rho_{T-1}}\right)^{-\alpha}\right] &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} d_T^{-\alpha}(u_{T-2}, v_{T-2}, \varepsilon_{T-1}, \eta_{T-1}). \\ &\quad \left(\frac{1+r}{1-\rho_{T-1}(u_{T-2}, v_{T-2}, \varepsilon_{T-1})}\right)^{-\alpha}\phi_{\varepsilon}(\varepsilon_{T-1})\,\phi_{\eta}(\eta_{T-1})\,d\varepsilon_{T-1}d\eta_{T-1} \end{aligned} \tag{21}$$

where $\phi_\varepsilon(\varepsilon_{T-1}) = \frac{1}{\sqrt{2\pi}} e^{-\frac{(\varepsilon_{T-1}-\sigma_\varepsilon/\sqrt{2})^2}{2}}$ and $\phi_\eta\left(\eta_{t-1}\right) = \frac{1}{\sqrt{2\pi}} e^{-\frac{(\eta_{T-1}-\sigma_\eta/\sqrt{2})^2}{2}}$ are the probability density functions of ε_{T-1} and η_{T-1} respectively. There is no analytical means of evaluating this double integral, and in such cases (where a normal distribution is concerned) the Gauss Hermite quadrature is commonly employed to obtain a numerical approximation.

We use a five-point Gaussian quadrature, which enables the expectation described by equation (21) to be approximated by evaluating the term inside the integral at 25 discrete combinations of the disturbance terms $\varepsilon_{T-1,i}$ and $\eta_{T-1,j}$. In fact, as $d_T = 1$ for all $\left(\varepsilon_{T-1}, \eta_{T-1}\right)$, the decision problem in period $T-1$ only requires 5 discrete evaluations (in ε_{T-1}, which determines ρ_{T-1}). After the expectation described by equation (21) is evaluated, the decision d_{T-1} at each grid intersection (u_{T-2}, v_{T-2}) is easily obtained.

Numerical approximations for the decisions in all periods, $65 \leq t \leq T-2$, are then obtained by recursive analysis as described above. Where the expectation $\left(\varepsilon_{t-1}, \eta_{t-1}\right)$ implies a combination (u_{t-1}, v_{t-1}) that is not located at a precise grid intersection of the period t state-space, a cubic spline is used to approximate the decision d_t.

There is one point worthy of note about the construction of the grids. Subsection 2.1 makes clear that the variances of u_t and v_t increase with time. There is greater uncertainty about the distant future than about the near future. In order to address this we adopt grids which expand with time. For u_t we consider nineteen points equally spaced covering the range $(-0.0293, 0.0293)(t-64)/46$. Since the variance of v_t rises in line with time we scale the grid points for v_t to $(t-64)^{\frac{1}{2}}$ using twenty-one[7] points in the range $(-0.0067, 0.0067)(t-64)^{\frac{1}{2}}$. Our results are not sensitive to the reducing the number of grid points, to nine and eleven respectively or, indeed to raising the number of grid points to thirty-nine and forty-one.

3 Results

We now present the dispersion of payments which would be made on pooled annuities for an investment of £1 by a 65 year old man in 2006, on the assumption that mortality rates follow the stochastic model described in section 2.1 and with a real interest rate of 2 1/2 per cent per annum. The results are based on a hundred thousand stochastic simulations of the model of mortality shown in equations (1 - 3). The graphs show the

---

[7] The use of different numbers of grid points for the two dimensions makes it easy to ensure that arrays are correctly defined in the use of the spline routines and is done for this reason.

payment in each year as a proportion of the initial investment made at age 65. The decile and quartile points shown, are those obtained by ranking the respective pay-outs in each period taken separately. A realisation of the stochastic processes which delivers a value on say the first decile in one year will probably not be on the first decile in other years.

Case 1: Moderate Risk Aversion $\alpha = 4$ Although a consensus has not yet emerged regarding a generally representative value for relative risk aversion, it is common for comparable studies to assume a value between 1.0 and 4.0.[8] Here we take the high end of this range to exaggerate the preference for insurance in our analysis. A more extreme preference for insurance is considered in Case 2, below. It can be seen in Figure 2 that the median pay-out rate subject to a relative risk aversion parameter of 4 rises to 6.40% by age 100. Its peak is 6.41% at age 104 and it falls back to 6.36% at age 110. This reflects the fact that payments are slightly front-loaded late in life because of the risk that the cohort will die completely before the age of 110. The 10th percentile pay-out at age 100 is 7.32% rising to 8.16% at age 110. The 90th percentile is 5.68% at age 100 falling to 5.43% at age 110. The mean pay-out rate (not shown) follows a path similar to that of the median until the age of 90 and then rises above it to a peak of 6.58% at age 107, falling to 6.54% by 110. Annuitants at age 65 are willing to accept what might seem a high degree of uncertainty about their dividends late in life because, of course, the chance of surviving to extreme old age is not very high; the decline in both the median and the mean in very old age reflects the fact that members of the pooled fund are concerned that all members will die before reaching the maximum age, so that some of the wealth in the fund will be wasted.

---

[8] Simulations undertaken by Low et al. (Forthcoming) are based upon a coefficient of risk aversion of 1.5, those by Auerbach & Kotlikoff (1987) assume 4, while Cooley & Prescott (1995) consider a value of 1. Grossman & Shiller (1981), Mankiw (1985) and Hall (1988) report econometric estimates for the intertemporal elasticity between 0 and 0.4, Blundell et al. (1994) report an estimate of 0.75, while Hansen & Singleton (1983) and Mankiw et al. (1985) report estimates just over 1. Values of the coefficient of risk aversion required to explain the equity premium puzzle (Mehra & Prescott (1985)) are large by comparison, although evidence from attitudinal surveys suggest that the value is unlikely to larger than 5 (Barsky et al. (1997)).

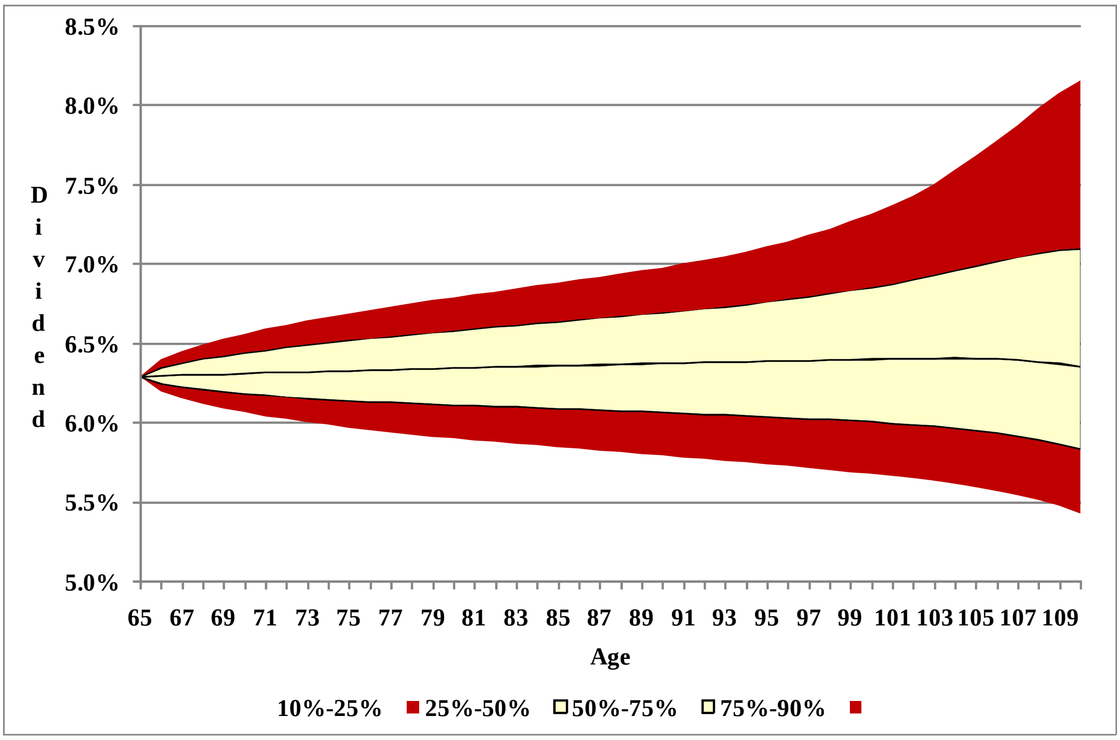

Figure 2: The Dispersion of Payouts from Pooled Annuity $\alpha = 4$

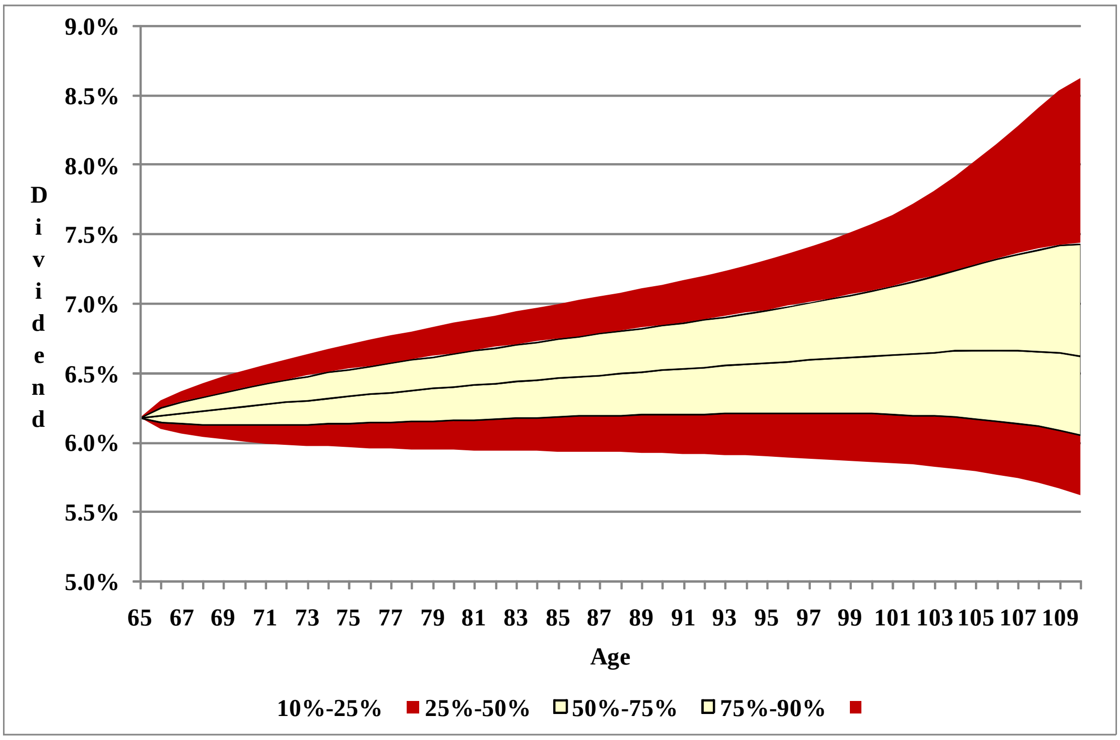

Figure 3: The Dispersion of Payouts from a Pooled Annuity $\alpha = 20$

Case 2: High Risk Aversion. $\alpha = 20$ Figure 3 shows the results for the second case, with $\alpha = 20$, i.e. with a highly risk-averse population. Comparison of figures 2 and 3 indicates clearly the impact of precautionary saving resulting from risk aversion. There is a more pronounced upward drift in the pay-out pattern shown in figure 3 than in figure 2. The initial pay-out rate with $\alpha = 20$ is reduced to 6.18% from 6.30% with $\alpha = 4$. The median pay-out rate then drifts up over time to reach 6.62% at age 100, and 6.63% at age 110 after a peak of 6.67% at age 106. At age 100 the lowest decile payment is 5.86% while the top decile payment is 7.57%. By age 110 the lowest decile is 5.62% while the top decile rises to 8.63%. The mean rises to a peak of 6.89% at age 108 , falling to 6.86% at age 110.

3.1 The Cost of Mortality Risk

In order to assess the amount that an annuitant might be prepared to pay to be protected from aggregate mortality risk, we need to assess first of all what would be the annuity rate paid by a risk-neutral seller of annuities, given the pattern of mortality shocks described in section 2.1. We note that, with any realisation of mortality rates, ρ_t the annuity rate which would leave the seller of the annuity exactly in balance is given, with

$\rho_0 = 0$ as

$$d = \left\{ \sum_{t=1}^{T} \frac{\prod_{\tau=0}^{t-1}(1-\rho_\tau)}{(1+r)^{t-1}} \right\}^{-1}$$

It follows that the risk-neutral seller of annuities sets

$$\hat{d} = E\left(\left\{ \sum_{t=1}^{T} \frac{\prod_{\tau=0}^{t-1}(1-\rho_\tau)}{(1+r)^{t-1}} \right\}^{-1} \right).$$

This can be evaluated by means of stochastic simulation over a large number of realisations of the ρ_t. We use fifty thousand realisations and employ the same realisations to evaluate the welfare of pooled annuitants and thus to identify the maximum that an annuitant would be willing to pay to be protected from aggregate mortality risk. This give an annuity rate of $\tilde{d}$ =6.33% p.a. as compared to the initial pay-out on the pooled annuity of 6.29% when $\alpha = 4$.

The initial level of welfare associated with an investment of w_1 is given as

$$\tilde{V}_1(w_1) = \sum_{t=1}^{T} \delta^{t-1} E\left\{ \prod_{\tau=1}^{t-1} (1-\rho_t)\, U\left(w_1\tilde{d}\right) \right\} = \frac{w_1^{1-\alpha}}{1-\alpha} \sum_{t=1}^{T} \delta^{t-1} E\left\{ \prod_{\tau=1}^{t-1} (1-\rho_t)\, \tilde{d}^{1-\alpha} \right\} \tag{22}$$

where the expectation is evaluated recursively.

For any given level of risk aversion, α we then find the value $\psi(\alpha)$ such that

$$V_1(w_{1,}\alpha) = \tilde{V}_1\{(1-\psi[\alpha])w_1\}$$

where $V_1(w_{1,}\alpha)$ is the value of $V(w_1)$ computed by equation (11) with parameter value α for the start of period 1 The homothetic nature of the problem means that this is independent of the value of w_1, as can be seen by comparing equation (22) with equation (11). Thus we find, noting that V_1 and $\tilde{V}_1$ are both negative

$$\psi(\alpha) = 1 - \left\{ \frac{\tilde{V}_1(1)}{V_1(1,\alpha)} \right\}^{\frac{1}{1-\alpha}}.$$

We calculate $\psi(\alpha)$ for $\alpha = 4$ and $\alpha = 20$ are derived from the simulations shown in more detail shown above. The values of $\psi(\alpha)$ are shown in table 1.

These results, 0.6% with standard risk aversion and 2.4% with very high risk aversion, set a maximum to the price for aggregate mortality risk which should be observed in the annuity market if annuitants can choose freely between conventional and pooled annuities. If sellers of annuities feel a need to charge more than this on conventional

$\boldsymbol{\alpha}$	4	20
$\psi(\alpha)$	0.6%	2.4%

Table 1: The Willingness to Pay for Protection from Aggregate Mortality Risk: Percentage of the Rate on a Conventional Annuity

annuities, then there should not be a market in the latter. They cannot be contrasted directly with studies such as Olivieri & Pitacco (2008*a*) because our framework does not allow us to say what the cost of carrying aggregate mortality risk is to an insurance company. In a market economy if retired people collectively want to shed aggregate mortality risk then that risk must be borne by young people as shareholders in life insurance companies and the amount that they charge for this depends on their tolerance of the risk[9]. It might well result in a market equilibrium where the market charge is lower than that shown in table However, the empirical evidence from the buy-out market discussed in the Introduction (Lane, Clark and Peacock 2008) points to a capital cost of 2.7% of the value of an annuity (and therefore the annual annuity payment) as being the charge for the capital needed to meet aggregate mortality risk imposed by an insurance company buying out a pension fund. Given this charge our analysis suggests that, unless the coefficient of risk aversion is much above 20 annuitants would prefer to carry aggregate mortality risk for themselves[10]. Since it is unlikely that people are as risk-averse as this, it is more probable that the market as it is forces people to be protected against aggregate mortality risk on terms which, given the choice, they would rather avoid.

4 Conclusions

It is possible, using the principles of a pooled annuity fund, to design variable annuities which protect sellers of annuities from aggregate mortality risk. Annuity payments are adjusted annually in a way which transfers the risk to annuitants in the light of emerging

---

[9] This issue is ignored in studies such as Friedberg & Webb (2006). Their analysis is based on the assumption that consumers are infinitely lived; indeed this is a key simplifying assumption of the consumption capital asset pricing model as it is conventionally applied, which raises obvious questions about its suitability to address aggregate mortality risk.

[10] There is a separate risk faced by annuity suppliers, that of adverse selection. But analysis using our model suggests that, unless this has an impact on the uncertainty surrounding the trend rate of improvement, the impact of this is small. Offsetting this, the buyout data relate to a typical portfolio of pensions. Such a portfolio obviously includes people older than 65 whose pensions are therefore less exposed to mortality risk. Practitioners have suggested that the capital charge imposed to address aggregate mortality risk on the pension of a 65-year old is of the order of 5%.

patterns of mortality. This obviously means that annuitants are exposed to the risks associated with variable income. If it is assumed that all income is consumed, it is possible to work out the extent allowance for risk aversion affects the desired pattern of annuity payments.

Simulations show that the impact of even quite high levels of risk aversion is relatively small. The explanation of this is that annuities are bought at the age of sixty-five. With a plausible model of mortality rates it is inevitable that the main uncertainty about mortality lies considerably in the future. For people in their sixties and early seventies, mortality rates are low in any case. Multiplicative shocks and trend changes have relatively little impact. Thus the main impacts of mortality uncertainty are substantially discounted by sixty-five year old annuitants and therefore they do not have much impact on the decisions of sixty-five year olds. We find that very high degrees of risk aversion are needed to justify the charge for aggregate mortality risk observed in the pensions buy-out market. If insurance companies selling annuities feel that they need to make large charges for carrying aggregate mortality risk, then financial regulators would do well to consider how to implement the alternative of making it possible for people to carry this risk for themselves. It is to be expected that this would result in a general increase in welfare.

This analysis has not discussed whether the degree of protection people require depends on the size of their pensions. Indeed the assumption of constant relative risk aversion assumed it does not. It is often suggested that people with relatively small pension funds are in particular need of protection from risk, with the implication that mortality-adjusted annuities might not be suitable for such people. However, the opposite may well be true. These calculations do not, of course, take any account of the effect on people's choices of the income they receive from state benefits, which are much more important for people with small annuities than for those with large ones. As Mitchell (2001) shows, the presence of a stable source of income in fact increases people's tolerance of the risk associated with other types of income. This suggests that people with modest privately-funded pensions have less need for protection from aggregate mortality risk than do those with larger retirement savings.

Of course insurance companies which are concerned that they face much larger aggregate mortality risk because they do not fully understand the pattern of the market to which they sell, might argue that these results present too optimistic a view of aggregate mortality risk. But the same approach can be used with different parameters to explore how mortality-adjusted annuities might operate in such circumstances and the gap between the charge in the pensions buy-out market and our estimate of willingness to pay with standard risk aversion suggests that the parameters would need to be very different indeed for annuitants to want to pay current market charges. In any case, the greater is the aggregate mortality risk faced by insurance companies selling annuities, the more important it is that they develop products which protect them from this risk. This paper sets out how to do that.

References

Ahlo, J. & Spencer, B. (1997), 'The Practical Specification of the Expected Error of Population Forecasts', *Journal of Official Statistics* **13**, 201–225.

Auerbach, A. J. & Kotlikoff, L. J. (1987), *Dynamic Fiscal Policy*, Cambridge University Press, Cambridge.

Barsky, R. B., Kimball, M. S., Juster, F. T. & Shapiro, M. D. (1997), 'Preference parameters and behavioral heterogeneity: an experimental approach in the health and retirement survey', *Quarterly Journal of Economics* **112**, 537–579.

Blundell, R., Browning, M. & Meghir, C. (1994), 'Consumer demand and the life cycle allocation of household expenditure', *Review of Economic Studies* **61**, 57–80.

Cooley, T. F. & Prescott, E. C. (1995), Economic growth and business cycles, *in* T. F. Cooley, ed., 'Frontiers of Business Cycles Research', Princeton University Press, Princeton.

Friedberg, L. & Webb, A. (2006), Life is Cheap: Using Mortality Bonds to Hedge Aggregate Mortality Risk. NBER Working Paper No 11984. www.nber.org/papers/W11984.

Grossman, S. J. & Shiller, R. J. (1981), 'The determinants of the variability of stock market prices', *American Economic Review* **71**, 222–227.

Hall, R. E. (1988), Intertemporal substitution in consumption. NBER Working Paper, W0720.

Hansen, P. & Singleton, K. (1983), 'Stochastic consumption, risk aversion and the temporal behaviour of asset returns', *Journal of Political Economy* **91**, 249–266.

Khalaf-Allah, M., Haberman, S. & Verall, R. (2006), ‘Measuring the Effects of Mortality Improvements on the Cost of Annuities’, *Insurance: Mathematics and Economics* **39**, 231–249.

Lane, Clark and Peacock (2008), Pension Buyouts . www.lcp.uk.com/information/documents/ LCPPensionBuyouts2008.pdf.

Lee, R. (1998), ‘Probabilistic Approaches to Population Forecasting’, *Population and Development Review* **24**, 156–190.

Leland, H. E. (1968), ‘Saving and Uncertainty: the Precautionary Demand for Saving’, *Quarterly Journal of Economics* **82**, 465–473.

Low, H., Meghir, C. & Pistaferri, L. (Forthcoming), ‘Wage risk and employment risk over the life cycle’, *American Economic Review* .

Mankiw, N. G. (1985), ‘Consumer durables and the real interest rate’, *Review of Economics and Statistics* **62**, 353–362.

Mankiw, N. G., Rotemberg, J. J. & Summers, L. H. (1985), ‘Intertemporal substitution in macroeconomics’, *Quarterly Journal of Economics* **100**, 225–251.

Mehra, R. & Prescott, E. C. (1985), ‘The equity premium: a puzzle’, *Journal of Monetary Economics* **15**, 145–161.

Mitchell, O. (2001), Developments in Decumulation: the Role of Annuity Products in Financing Retirement. NBER Working Paper No 8567.

Olivieri, A. & Pitacco, E. (2008*a*), ‘Assessing the Cost of Capital for Longevity Risk’, *Insurance: Mathematics and Economics* **42**, 1013–1021.

Olivieri, A. & Pitacco, E. (2008*b*), Stochastic Mortality: the Impact on Target Capital. Carefin Working Paper No 15/08 Universita Bocconi.

Pensions Commission (2005), Second Report: Appendix E. Uncertainties in Life Expecantcy Projections. www.pensionscommission.org.uk/publications/2005/ annrep/appendices/appendix-e.pdf.

Piggott, J., Valdez, E. & Detzel, B. (2005), ‘The Simple Analytics of a Pooled Annuity Fund’, *Journal of Risk and Insurance* **72**, 497–520.

Pratt, J. (1964), 'Risk Aversion in the Small and in the Large', *Econometrica* **32**, 122–136.

Renshaw, A., Haberman, S. & Hatzopoulos, P. (1996), 'Modelling of Recent Mortality Trends in UK Male Assured Lives', *British Actuarial Journal* **5**, 96–124.

Stamos, M. (2008), 'Optimal Consumption and Portfolio Choice for Pooled Annuity Funds', *Insurance: Mathematics and Economics* **43**, 56–68.