Risk and Mortality-adjusted Annuities

Justin van de Ven and Martin Weale*

National Institute of Economic and Social Research
2, Dean Trench Street,
London SW1P 3HE

National Institute Disucssion Paper No 322
29th August 2008

Abstract

This paper discusses the way in which payments from pooled annuity funds need to be adjusted to take account of the fact that future mortality rates are uncertain. Mortality-adjusted annuities, as we describe payments from the pooled fund are variable annuities in which aggregate mortality risk is transferred from the seller of annuitees to the annuitants. If annuitants are risk averse the payments from the fund should be adjusted to reflect this. We show how the adjustment can be calculated and compare the payment profiles from a risk-adjusted funds with alternatives which either ignore uncertainty completely or take account of the uncertainty but assume that annuitants are not risk averse. It is shown that, even for very risk-averse annuitants, initial payments are reduced only slightly to provide acceptable insurance against the implications of uncertainty about future mortality rates.

JEL: D14, D91, J11, J14
Keywords: Variable Annuity, Aggregate Mortality Risk, Risk Aversion

*We gratefully acknowledge support from the ESRC. An earlier version of this paper was presented at the FIRM conference in Manchester. June 2008.
Contact address mweale@niesr.ac.uk.

1 Introduction

The arithmetic of annuity calculations when mortality rates are certain has been known for the best part of two centuries. However, the problems arising from uncertain mortality rates are only now facing attention. Ahlo & Spencer (1997), Lee (1998) and Renshaw, Haberman & Hatzopoulos (1996) have addressed the problem of producing 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- or aggregate mortality risk- for the insurance companies selling annuities. Khalaf-Allah, Haberman & Verall (2006) look at the distribution of the annuity cost face by an insurance company selling them. Olivieri & Pitacco (2008) investigate the implications of aggregate mortality risk for the cost of capital faced by an insurance company.

Piggott, Valdez & Detzel (2005) discuss the operation of a pooled annuity fund- one in which a fund is set up in which the assets of the decedents accrue to the survivors. Such a fund provides protection from individual mortality risk, but leaves annuitants' incomes and consumption possibilities subject to the uncertainty associated with the mortality risk of the group. It offers a means by which sellers of annuities can protect themselves against the effects of uncertainty about aggregate mortality risk and, as they point out, offers an alternative to the "survivor bonds" proposed by Blake & Burrows (2001).

The anlaysis presented by Piggott *et al.* explores how annuity payments may be varied in the light of deterministic, but unanticipated, mortality shocks, and examines how the payments might vary as a consequence of stationary random shocks to mortality rates. The focus of this article is on the management of the payments from a pooled annuity fund when mortality both mortality rates and trends in mortality rates are subject to persistent random disturbances. We show that annuitants who are concerned about the risks to future annuity payments will want the fund to pay out, in its early stages, less than it would if they were indifferent to those risks.

We use dynamic programming methods to evaluate the optimal pay-out path of a mortality adjusted annuity. This optimal path depends, of course, on the assumed attitude of the annuitant to risk, as well as to the assumptions made about uncertainty in the mortality rate. An annuitant who is very concerned about risk will choose a low initial pay-out so that, if mortality rates prove much lower than expected, the annuity payment can nevertheless be maintained. Should mortality rates meet expectations this means that the annuity payment will rise over time. 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 the annuity payment will be cut if mortality rates fail to meet expectations.

We illustrate this by exploring three cases. In the first case the payment from the pooled fund is recalculated each year, but, as Piggott et al. (2005) proposed, on the basis of point forecasts of future mortality rates. Thus the possible dispersion of future mortality rates is ignored and no allowance is made for risk aversion. In the second case the payment is calculated so as to maintain the expected payment constant in the light of the dispersion of possible future mortality rates, but no allowance is made for risk aversion on the part of the annuitants. In the third case the payments are calculated on the assumption that annuitants are risk averse and, in effect, wish to insure themselves against future uncertainty.

Section 2 describes briefly the way in which uncertainty about future mortality rates complicates the calculation of the payment from a conventional annuity. In section 3 we present the problem faced in calculating the optimal consumption path of the holder of a mortality-adjusted annuity. Section 4 provides account of the structure we assume for shocks to mortality. In section 5 we specify the three cases in greater detail and in section 6 explain how we calculate numerical solutions for the annuity payments in each case. In section 7 we provide estimates of the pay-out path for a male sixty-five year old annuitant in the United Kingdom, in each of the three cases. Section 8 draws conclusions.

2 Uncertain Mortality

The problem we address is well-known and needs only the briefest explanation. We consider an annuity designed to allow the annuitant to maintain a constant standard of living; we define everything in real terms and work with the real rate of interest. The real rate of interest is taken to be a constant[1], r. For an annuity costing £1, which provides a constant payment at the start each year, the payment for someone of age t is, in the absence of uncertainty, with a mortality rate for the relevant cohort of ρ_ι in year ι and with a maximum life-span of T,

$$d_t = \frac{1}{1 + \sum_{\tau=t}^{T-1} \prod_{\iota=t}^{\tau} \frac{1-\rho_\iota}{1+r}} \tag{1}$$

The difficulty is that, at the start of year t none of the subsequent mortality rates is known with certainty. Thus the conventional annuity is a risky proposition for an insurance company to sell. It can compensate for that risk by offering a payment below d_t, but one would expect

[1] Rates of return may also be uncertain in which case an approach similar to that set out here can be used to address the issue. But our focus is on uncertain mortality rates.

there also to be a market for pooled-fund annuities in which annuitants rather than insurance companies carry the risk associated with the uncertainty of ρ_ι Even if the annuity is not limited to a pooled fund, insurance companies may well prefer to issue payments which depend on the evolving mortality rates of a complete cohort; this would protect them from the risks associated with general uncertainty about future mortality rates. In this paper we explore how such a variable annuity, which we describe as a mortality-adjusted annuity[2] should address the risks associated with uncertain mortality rates.

3 The Optimisation Problem

The problem we address has its roots in the analysis of consumer behaviour under uncertainty. In the absence of uncertainty, economic theory suggests that the time profile of consumption that consumers choose depends on their subjective discount rates 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. If the reverse is true consumption should be expected to fall over time. When they are equal consumers plan to keep consumption constant, of course revising their consumption in the face of unexpected shocks. Leland (1968) shows that a 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, as people's resources increase, they become more tolerant of any given absolute degree of uncertainty about that future income. The utility function we specify,

$$U(c_t) = \frac{c_t^{1-\alpha}}{1-\alpha} \quad , \tag{2}$$

is standard in the analysis of problems of this type. It assumes constant relative risk aversion, so that it indeed has the property that, as people's wealth increases, their tolerance of any absolute degree of uncertainty about future income diminshes[3]. Constant relative risk aversion implies that the effect of proportionate uncertainty about future income is independent of resources so that the impact of £1 of uncertainty is bound to decrease as initial wealth rises.

The presence of precuationary saving is highly relevant to our problem because an annuity has the effect of sharing the assets of those who die in each year among the survivors, thus raising the return on their capital. Given our specification of the utility function, uncertainty about future mortality rates means, therefore, that, consumers who would aspire to a constant level of real consumption in the absence of aggregate mortality risk will instead reduce their

[2] We assume that the annuity pays an annual payment at the start of each year to annuitants alive then. If payments were made more frequently, say monthly, but information on mortality accrued only annually, our problem would be more complicated but the same general principles would apply.

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

consumption in the early years of the annuity so that a reserve is available should the mortality rate of their cohort of annuitants be lower than expected. This means that, on average they will want a path with annuity payments which start low and rise over time. The extent to which they want to depart from the level path will depend on their attitude to risk.

We assume that the mortality-adjusted annuity pays out the amount that the annuitant would rationally choose to consume in each period. This means that any precautionary saving which is desirable takes place inside the fund, and allows the benefits of such savings to be protected from individual longevity risk. In other words, by setting pay-outs on to the consumption path, the full benefits of annuitisation are delivered, subject to the annuitant carrying the aggregate mortality risk.This would not be the case if annuitants were left to make their own arrangements for handling the problem of uncertainty about future aggregate mortality rates.

We consider a fund of value w_t at the start of the year in which the cohort is aged t. The annuitant wishes to choose an optimal consumption stream given that the mortality rate associated with the cohort to which the annuitant belongs is uncertain. If ρ_t is the mortality rate of the cohort in year t then the optimisation problem can be written as a Bellman equation

$$V\left(w_t\right)=\underset{c_t}{Max}\,U\left(c_t\right)+\delta E\left(V\left(w_{t+1}\right)\right) \tag{3}$$

where δ is an exogenous discount factor and

$$w_{t+1}=\frac{1+r}{1-\rho_t}\left(w_t-c_t\right) \tag{4}$$

since it is assumed that the remaining assets of the members of the cohort 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 as too, therefore, are future consumption possibilities.

$V(w_t)$ is the total remaining expected life-time welfare as a function of w_t, on the assumption that optimal consumption choices are made. $V\left(w_{t+1}\right)$ represents the welfare derived from wealth conditional on being alive at the start of year $t+1$ with $V\left(w_{t+1}\right)=0$ in the event of death. Since the probability of surviving into period $t+1$ is, of course, 1-ρ_t, equation (3) becomes

$$V\left(w_t\right)=\underset{c_t}{Max}\,U\left(c_t\right)+\delta E\left(1-\rho_t\right)\left\{V\left(\frac{1+r}{1-\rho_t}\left[w_t-c_t\right]\right)\right\} \tag{5}$$

In the final period T which represents the maximum possible life-span

$$V(w_T)=\underset{c_T}{Max}\,U(c_T) \qquad c_T\leq w_T \tag{6}$$

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

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

$$V\left(w_{T-1}\right)=\underset{c_{T-1}}{Max}\,U\left(c_{T-1}\right)+\delta E\left(1-\rho_{T-1}\right)\left\{U\left(\frac{1+r}{1-\rho_{T-1}}\left[w_{T-1}-c_{T-1}\right]\right)\right\} \tag{7}$$

With $U(c_t)=\frac{c_t^{1-\alpha}}{1-\alpha}$ and $c_{T-1}=d_{T-1}w_{T-1}$ where d_{T-1} is the payment in period $T-1$ on annuitised capital of £1 at the start of period $T-1$, then

$$V\left(w_{T-1}\right)=w_{T-1}^{1-\alpha}\left(\underset{d_{T-1}}{Max}\,\frac{d_{T-1}^{1-\alpha}}{1-\alpha}+\frac{\delta}{1-\alpha}E\left\{\left(1-\rho_{T-1}\right)\left(\frac{1+r}{1-\rho_{T-1}}\right)^{1-\alpha}\left(1-d_{T-1}\right)^{1-\alpha}\right\}\right) \tag{8}$$

In the general case for $65\leq t\leq T-1$, where $t=1$ refers to the age at which the annuity is transacted, taking advantage of the structure of the utility function and since

$$w_{t+1}=w_t\left(1-d_t\right)\left(\frac{1+r}{1-\rho_t}\right) \tag{9}$$

$$V\left(w_t\right)=w_t^{1-\alpha}\left(\underset{d_t}{Max}\,\frac{d_t^{1-\alpha}}{1-\alpha}+\frac{\delta}{1-\alpha}E\left\{V\left(\left\{1-d_t\right\}\left\{\frac{1+r}{1-\rho_t}\right\}\right)\right\}\right). \tag{10}$$

The solution to (10) is found from the first-order conditions represented by the Euler equation

$$\begin{aligned}d_t^{-\alpha}&=\delta(1+r)E\left\{\left(1-d_t\right)^{-\alpha}d_{t+1}^{-\alpha}\left(\frac{1+r}{1-\rho_t}\right)^{-\alpha}\right\}\\&=\delta(1+r)\left(1-d_t\right)^{-\alpha}E\left\{d_{t+1}^{-\alpha}\left(\frac{1+r}{1-\rho_t}\right)^{-\alpha}\right\}\end{aligned} \tag{11}$$

We set $\delta=1/(1+r)$ so that future utility is discounted at the rate of interest, and, in the absence of uncertainty, the desired level of consumption is constant over time. This gives

$$d_t^{-\alpha}=E\left\{d_{t+1}^{-\alpha}\left(\frac{1+r}{1-\rho_t}\right)^{-\alpha}\right\}\left(1-d_t\right)^{-\alpha} \tag{12}$$

so that

$$d_t=\frac{1}{1+\left(E\left\{d_{t+1}^{-\alpha}\left(\frac{1+r}{1-\rho_t}\right)^{-\alpha}\right\}\right)^{\frac{1}{\alpha}}} \tag{13}$$

This recursive equation provides the basis for evaluating the optimal annuity rate in each period. In order to apply equation (13) it is necessary to adopt some specification for the nature of the disturbances to the mortality rate. We do this in the next section.

4 Mortality Disturbances

We adopt a model of mortality disturbances which focuses on the uncertainty surrounding the trend rate of decline in log mortality. It is simpler than some of the approaches referred to in the introduction. In part this is because, since when looking at annuities, we are concerned only with the mortality rates of individual cohorts, 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 follow the process

$$\log \rho_t = \log \rho_t^* + u_t \text{ if } \log \rho_t^* + u_t < \log(0.8) \text{ for all } t \geq 65 \quad (14)$$

$$\log \rho_t = \log(0.8) \text{ if } \log \rho_t^* + u_t \geq \log(0.8) \text{ for all } t \geq 65 \quad (15)$$

$$\log \rho_t = \log(0.8) \text{ if } \log \rho_{t-1} = \log(0.8) \quad (16)$$

$$u_t = u_{t-1} + \theta_t + \varepsilon_t + v_{t-1} \quad (17)$$

$$v_t = v_{t-1} + \eta_t \quad (18)$$

$$u_{64} = v_{64} = 0 \quad (19)$$

Here $\varepsilon_t \tilde{} N(-\frac{\sigma_\varepsilon^2}{2}, \sigma_\varepsilon^2)$ and so too is $\eta_t \tilde{} N(-\frac{\sigma_\eta^2}{2}, \sigma_\eta^2)$. These definitions ensure that $E(e^{\varepsilon_t}) = E(e^{\eta_t}) = 1$

Equation (14) says that log mortality is equal to some reference value plus a random term, u_t. However with no constraint there is a risk that the mortality rate generated by the model might rise above one. We impose in equation (15) a maximum rate of 0.8 and also, in equation (16) that if, at some age, mortality rises to 0.8, it does not then fall back below it.

In using a model of this type it is necessary to relate the mortality rates to those provided by some exogenous source such as the UK Government Actuary whose figures we use here. While it is often unclear whether point projections are in fact expectations (rather than for example medians), we assume that the user of the model wants to calibrate it so that $E(\rho_t) = \tilde{\rho}_t$ where $\tilde{\rho}_t$ is a series of exogenously given cohort mortality rates. There are two distinct effects to be addressed and we consider them separately. First of all we want to ensure that $E(e^{u_t}) = 1$ and secondly that with $E(e^{u_t}) = 1, E(\rho_t) = \tilde{\rho}_t$. ρ_t^* is the sequence of time varying non-stochastic terms which ensures that this is true. It should be noted that, in the absence of the trunctation implied by equations 14 and 15 $\rho_t^* = \tilde{\rho}_t$.

The specification of η_t implies that $Var(v_t)$=$(t-64)\sigma_\eta^2$. Since, with this definition of v_t, $E(v_t) = -(t-64)\sigma_\eta^2/2$, we also have that $E(e^{v_t}) = 1$. However, u_t requires the extra term θ_t if we are to ensure that $E(e^{u_t}) = 1$ as is required if we are to offset the drift arising from the second-order process.This term is computed as follows. We have

$$u_t = \sum_{i=65}^{t} \varepsilon_i + \sum_{i=65}^{t} (t-i)\,\eta_i \tag{20}$$

Thus

$$E(u_t) = E(u_{t-1}) + \theta_t - \frac{\sigma_\varepsilon^2}{2} - (t-65)\frac{\sigma_\eta^2}{2}; Var(u_t) = (t-64)\sigma_\varepsilon^2 + \sigma_\eta^2 \sum_{i=1}^{t} (t-i)^2 \quad \text{and} \tag{21}$$

$$Var(u_t) = Var(u_{t-1}) + \sigma_\varepsilon^2 + (t-65)^2\sigma_\eta^2. \tag{22}$$

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

$$\theta_t = \frac{\sigma_\eta^2 \left\{(t-65) - (t-65)^2\right\}}{2}. \tag{23}$$

With $E(e^{u_t}) = 1$ we still have to address the fact that the distribution of $\log \rho_t$ is truncated. This cannot be done analytically. We use the numerical procedure described in section 7 to find a sequence of ρ_t^* such that $E(\rho_t) = \tilde{\rho}_t$

5 Three Approaches to Annuity Management

The specification of the mortality process above allows us to specify in more detail the three cases mentioned in the introduction.

5.1 Case 1: No allowance for risk and uncertainty.

In this case we assume that the annuity is managed in a manner which ignores the uncertainty surrounding future mortality rates. Point forecasts of mortality rates are produced. The point forecast for period τ based on observations up to the start of period t is given, using the most recent information on the shock to mortality, u_{t-1} and the most recent value of the trend v_{t-1} as

$$\log \hat{\rho}_\tau(u_{t-1}, v_{t-1}) = \log \rho^*_{T-1} + u_{t-1} + (\tau + 1 - t)v_{t-1}. \tag{24}$$

It is straightforward to see that the payment, $\hat{d}_t(u_{t-1}, v_{t-1})$, as a proportion of the value of the fund at the start of period t, is given as

$$\hat{d}_t(u_{t-1}, v_{t-1}) = \frac{1}{1 + \sum_{\tau=t}^{T-1} \prod_{\iota=t}^{\tau} \frac{1-\hat{\rho}_\iota(u_{\iota-1}, v_{\iota-1})}{1+r}}. \tag{25}$$

The payment will be revised each year in the light of new information on the mortality shock and the trend rate of change.

5.2 Case 2: Allowance for uncertainty. $\alpha = 0$

The second case of interest arises when the fund is run in a manner which takes account of the effect of uncertainty on future mortality rates but where annuitants are assumed to be risk neutral. This is represented with $\alpha = 0$ and implies that the fund does not undertake precautionary saving.

We can solve this problem recursively in a manner similar to that employed in section 4. Since it is not clear how to identify the limiting case of equation (13) as $\alpha \to 0$, we use work from equation (26) below.

$$\tilde{d}_{t-1}(u_{t-2}, v_{t-2}) = \{1 - \tilde{d}_{t-1}(u_{t-2}, v_{t-2})\} \frac{(1+r)\tilde{d}_t(u_{t-1}, v_{t-1})}{1 - \rho_{t-1}(u_{t-2}, v_{t-2})} \tag{26}$$

This sets the payment out of an initial fund of £1 in year t equal to the payment in year $t+1$, which can be paid out of the capital remaining from the distribution in year t after interest has been earned on it and after adjusting for the mortality of the relevant population in year t. Taking expectations

$$\frac{\tilde{d}_t(u_{t-1}, v_{t-1})}{1 - \tilde{d}_t(u_{t-1}, v_{t-1})} = E\left\{ \frac{(1+r)\tilde{d}_{t+1}(u_t, v_t)}{1 - \rho_t(u_{t-1}, v_{t-1})} \right\} \tag{27}$$

This gives, if we want to keep the expected payment constant

$$\tilde{d}_t(u_{t-1}, v_{t-1}) = \frac{E\left\{ \frac{(1+r)\tilde{d}_{t+1}(u_t, v_t)}{1-\rho_t(u_{t-1}, v_{t-1})} \right\}}{1 + E\left\{ \frac{(1+r)\tilde{d}_{t+1}(u_t, v_t)}{1-\rho_t(u_{t-1}, v_{t-1})} \right\}} \tag{28}$$

Equation (28) can be solved in the same way as equation 13, by means of the Gaussian quadrature as described in section 6 below.

5.3 Case 3: Allowance for risk and uncertainty. $\alpha = 20$

The third case we examine is that set out by equation 13. We use a value of $\alpha = 20$. This is believed large compared to most studies of people's attitude to risk, which suggest values of $\alpha = 1$ to $\alpha = 5$. We have adopted a large value in order to produce results which are clearly distinguishable from those of case 2 and to set a limit to reasonable allowance for risk. We now discuss in detail the solution method for cases 2 and 3 in detail.

6 Model Solution

We work backward from the start of period $T-1$. At this point it is necessary to forecast $\log \rho_{T-1}$ using the values of u_{T-2} and v_{T-2} which are assumed to be known.

$$\log \rho_{T-1} = \log \rho^*_{T-1} + u_{T-2} + \theta_{T-1} + v_{T-2} + \eta_{T-1} + \varepsilon_{T-1} \tag{29}$$

Since the error process has two disturbance terms, η_{T-1} and ε_{T-1} it is necessary to integrate over two dimensions. The complications arise over evaluating the expectational term in the denominator of equation (13).

$$E\left\{ d_T^{-\alpha} \left(\frac{1+r}{1-\rho_{T-1}} \right)^{-\alpha} \right\} \tag{30}$$

$$= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} d_T^{-\alpha}(u_{T-2}, v_{T-2}, \varepsilon_{T-1}, \eta_{T-1})$$
$$\left(\frac{1+r}{1-\rho_{T-1}(u_{T-2}, v_{T-2}, \varepsilon_{T-1})} \right)^{-\alpha} \phi_\varepsilon(\varepsilon_{T-1})\, \phi_\eta(\eta_{T-1})\, d\varepsilon_{T-1} d\eta_{T-1} \tag{31}$$

where $\phi_\varepsilon(\varepsilon_{T-1}) = \frac{1}{\sqrt{2\pi}} e^{-\frac{(\varepsilon_{T-1}-\sigma_\varepsilon/\sqrt{2})^2}{2}}$ and $\phi_\eta(\eta_{t-1}) = \frac{1}{\sqrt{2\pi}} e^{-\frac{(\eta_{T-1}-\sigma_\eta/\sqrt{2})^2}{2}}$ are the probability density functions of ε_{T-1} and η_{T-1} respectively. There is no analytical means of evaluating this double integral. However, Gaussian quadrature offers a method of numerical evaluation.

We use five-point Gaussian quadrature which allows us to evaluate the double integral by providing weights, π_i and π_j associated with five specified values of the disturbances $\varepsilon_{T-1,i}$ and $\eta_{T-1,j}$ for given values of u_{T-2} and v_{T-2}; these weights are chosen to be optimal given that the underlying variables are normally distributed. We define $\rho_{T-1,i,j}$ as

$$\log \rho_{T-1,i,j} = \log \rho^*_{T-1} + u_{T-2} + v_{T-2} + \theta_{T-1} + \varepsilon_{T-1,i} + \eta_{T-1,j} \tag{32}$$

and, for any value of $t < T$

$$u_{t,i} = u_{t-1} + \theta_t + v_{t-1} + \varepsilon_{t,i} \tag{33}$$
$$v_{t,j} = v_{t-1} + \eta_{t,j}.$$

We then evaluate

$$d_{T-1}(u_{T-2}, v_{T-2}) = \frac{1}{1 + \left(\sum_{i,j} \pi_i \pi_j \left\{ d_T^{-\alpha}(u_{T-2}, v_{T-2}, \varepsilon_{T-1,i}, \eta_{T-1,j}) \left(\frac{1+r}{1-\rho_{T-1}(u_{T-2}, v_{T-2}, \varepsilon_{T-1,i})} \right)^{-\alpha} \right\} \right)^{\frac{1}{\alpha}}} \tag{34}$$

at the specified values of u_{T-2} and v_{T-2}. This allows us to construct a grid of values of d_{T-1} as functions of u_{T-2} and v_{T-2}. We can then roll equation (34) back one period in order to evaluate d_{T-2}. In order to do this we have to interpolate and extrapolate the grid of values for

d_{T-1} in order to find the values at the required quadrature points. We do that using a cubic spline. Repeating this recursively we can produce grids for $d_t(u_{t-1}, v_{t-1})$ for all $T > t \geq 65$, and thereby establish propensities to consume out of wealth at all ages as functions of known values of the disturbances.

There is one point worthy of note about the construction of the grids. Section 4 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[4] 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.

7 An Illustration

7.1 The Dispersion of Mortality Rates

We illustrate the workings of risk-adjusted annuities with reference sixty-five year old men, using the mortality rates which underpin the UK's official cohort life expectancy projections for 2006. These mortality rates were kindly provided by the Government Actuary. We assume that these mortality rates represent expected mortality rates.

After a certain amount of experimentation we set $\sigma_\varepsilon = 0.01$ and $\sigma_\eta = 0.03$. These parameters have to be set before it is possible to determine the sequence ρ_t^*. If our model of stochastic mortality were simply the log- normal model of section 4 the mean values of the disturbances would offset the drift of the log-normal process and the model, applied to the mortality rates $\tilde{\rho}_t$ provided by the Government Actuary for men aged sixty-five in 2005 would have the property that these were also the expected mortality rates[5]. However, the upper limit to the mortality rate defined in equation (15) implies that the mean mortality rates generated by simulating the model are lower than $\tilde{\rho}_t$.

We resolve this problem iteratively. Starting with the values of $\tilde{\rho}_t$ provided by the Government Actuary, we set $\rho_t^{*0} = \tilde{\rho}_t$ as an initial estimate of ρ_t^*.We then simulate the stochastic model of mortality rates given by equations (14)- (18). The mean values of the mortality rates resulting from one hundred and fifty thousand simulations, are denoted $\rho_t^{0,1}$ We set $\rho_t^{*1} = \rho_t^{*0}.\rho_t^{*0}/\rho_t^{0,1}$. We repeat this process with $\rho_t^{*k} = \rho_t^{*k-1}.\rho_t^0/\rho_t^{k-1,k}$. It can be seen that ρ_t^{*k} stabilises as $\rho_t^0/\rho_t^{k-1,k}$ converges to 1, i.e. when the use of ρ_t^{*k} as inputs to the model delivers the required expected

[4]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.

[5]The programming was verified by means of one hundred and fifty thousand simulations.

mortality rates of $\tilde{\rho}_t$. We carry out ten iterations. We then use $\rho_t^* = \rho_t^{*10}$ to provide the mortality rates which drive the stochastic model since by this stage further iteration did not seem to offer visible improvement[6]. 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 0.1 percentage points.

These values of ρ_t^*, σ_ε and σ_η give a mean life expectancy[7] at the age of sixty-five of 20.6 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 a 90% confidence interval of 18.9 to 22.6 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 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. 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.

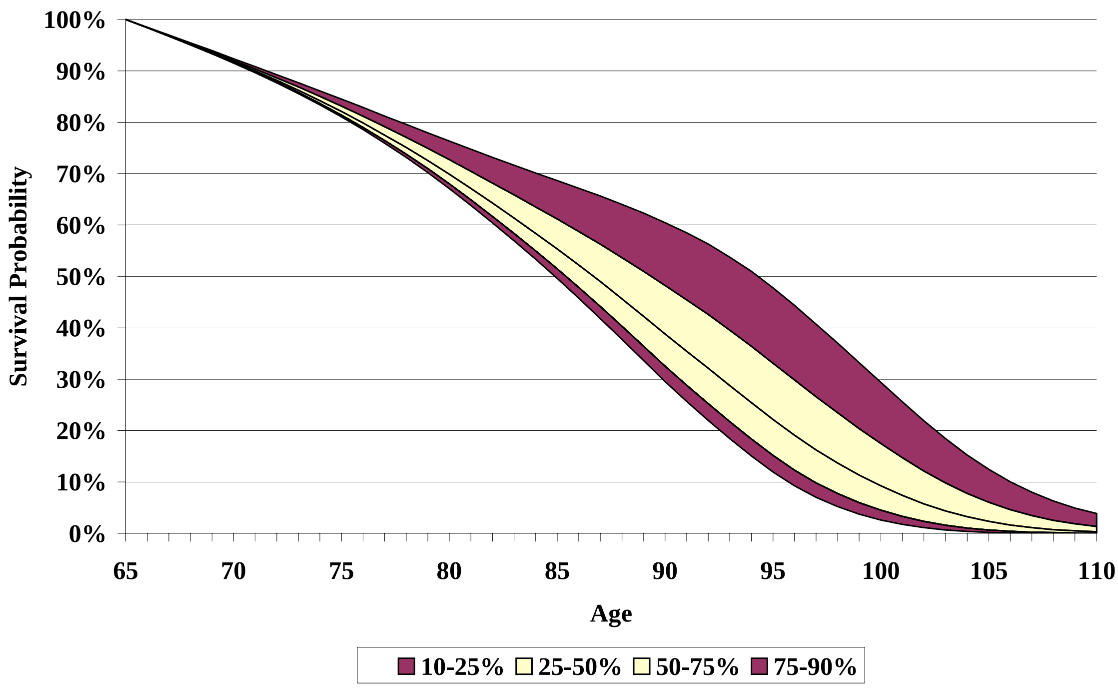

Figure 1: The Distribution of Survival Rates by Age of an Initial Population of Sixty-five Year Old Men

[6]It is not possibel to apply a standard convergence test because, even with one hundred and fifty thousand simulations there is an element of stochastic variation present.

[7]Based on fifty thousand simulations.

7.2 Annuity Rates by Age

We define the annuity rate at any age as the annual annuity payment which would be made from a sum of £1 invested on reaching that age. We show first in figure 2 the percentage point differences in the annuity rates per pound of annuitised capital as a function of the age of the annuitant for the cases 2 and 3 measured relative to case 1. In all cases these are the payment made when the realised values of u_t and v_t are zero. So with $\alpha = 0$ and $\alpha = 20$ the funds are managed to allow for shocks which do not in fact materialise. The outcomes are measured relative to the annuity rate which would be paid if mortality rates took values $\tilde{\rho}_t$ and there were no uncertainty.

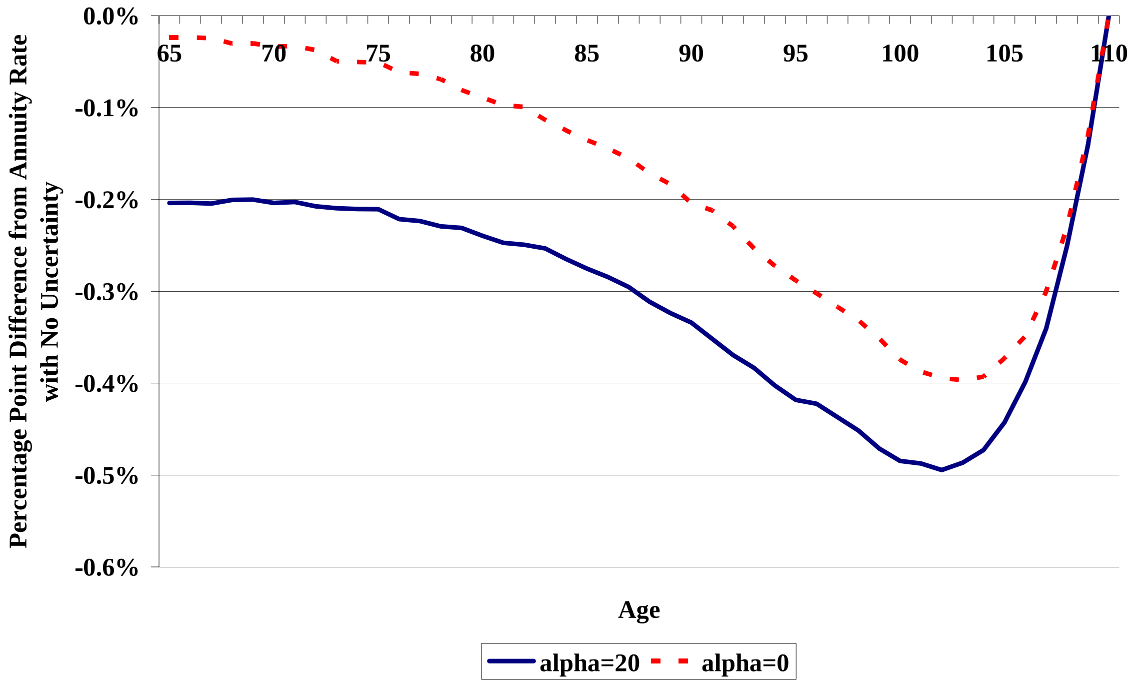

Figure 2: Deviations of Annuity Rates from those Paid in the Absence of Uncertainty

For a sixty-five year old man the rate is depressed by 0.02 percentage points in the absence of risk aversion ($\alpha = 0$). With the very high level of risk aversion implied by $\alpha = 20$, the annuity rate is still depressed by only 0.2 percentage points. As age increases, up to age one hundred and two the gap relative to no uncertainty increases. But the fact that the curve with $\alpha = 20$ become closer to that with $\alpha = 0$ mean that this is increasingly because of the dispersion in mortality rates rather than because of risk aversion. A reasonable, and not very surprising conclusion, is that the older people are the less mortality risk their cohort faces because there is less time for factors such as uncertainty about trend rates of change of mortality to have an effect. These results imply that, if insurance companies feel they need to offer much lower rates to protect themselves from mortality risk, then welfare would probably be increased if instead annuitants were given the opportunity to carry this risk for themselves.

7.3 Outcomes

We now present the dispersion of payments which would be made on variable annuities for an investment of £1 by a sixty-five year old man on the assumption that mortality rates follow the stochastic model described in section 4. The results are based on fifty thousand stochastic simulations of the model of mortality shown in equations (14 - 18). The graphs show the payment in each year as a proportion of the initial investment made at age sixty-five. The decile and quartile points show are those 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: No allowance for risk and uncertainty. In the first case, shown in figure 3 we show the dispersion of payments if they are revised in each year as a result of updated information on the trend and level of mortality rates, but if no account is taken of the stochastic nature of disturbances to mortality. Thus the payments are calculated using formula (25) with the term θ_t in equation (17) being ignored. It can be seen that in this case median payments drift down, with the drift becoming increasingly marked with age. This drift arises because an actuary managing a fund on this basis is, in ignoring future disturbances, taking no account of the non-linear way in which forecast future mortality rates enter the equation (1).

Case 2: Allowance for uncertainty. $\alpha = 0$ The second case arises where the payments are calculated in a manner which takes full account of the dispersion of future mortality rates but in which no allowance is made for the risk aversion of annuitants. It can be seen that the median is now stable until age 100, after which it drifts down very slightly; the mean (not shown) by contrast rises from 6.33% to 6.34%- a movement attributable to sampling variation. If annuitants are assumed risk-neutral the initial payment should be set so that the expected payment is constant.

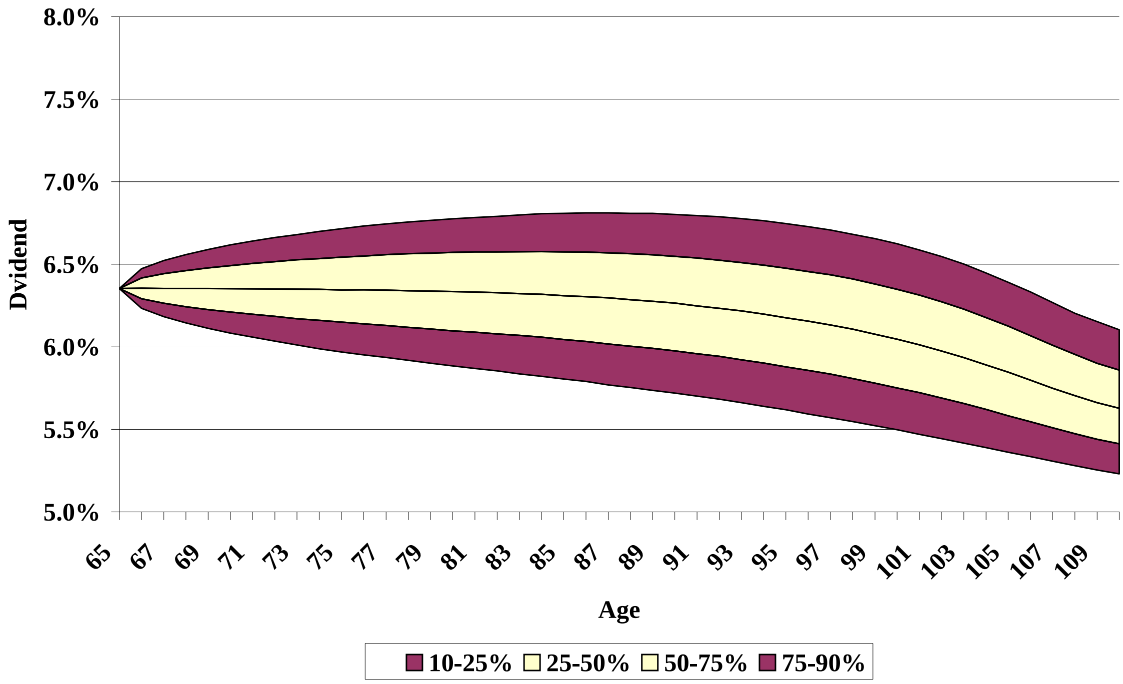

Figure 3: Case 1: Payment Rates on a Mortality-adjusted Annuity which takes no Account of the Dispersion of Mortality Rates

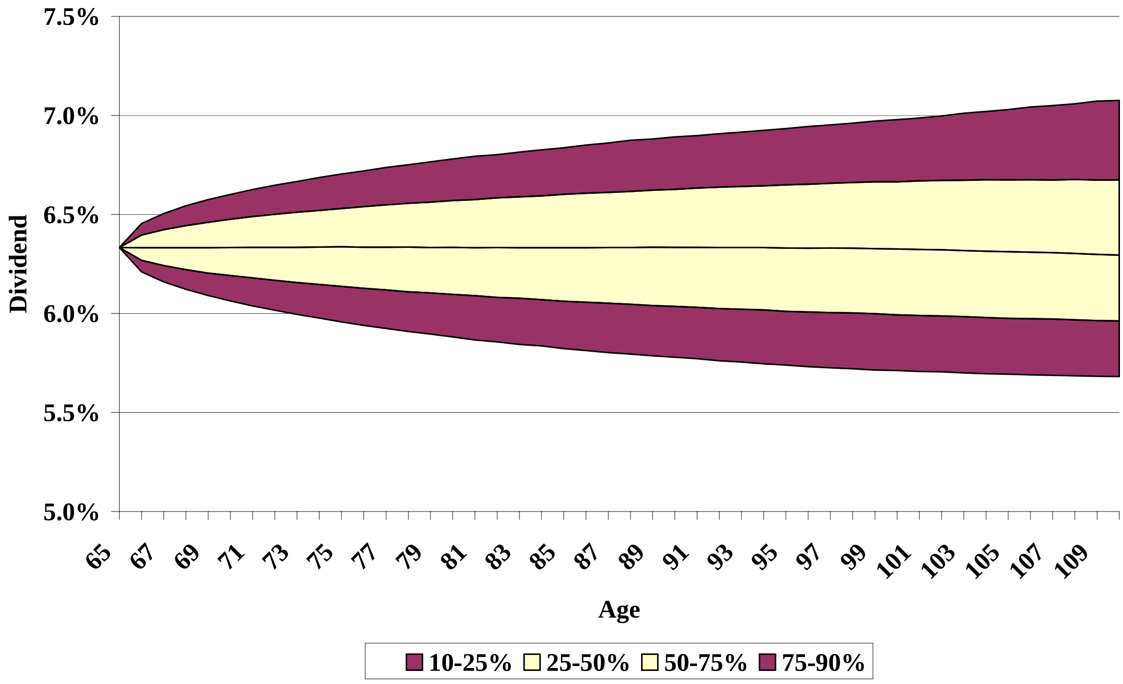

Figure 4: Case 2: Payment Rates on a Mortality-adjusted Annuity which takes account of the Dispersion of Mortality Rates but with $\alpha = 0$

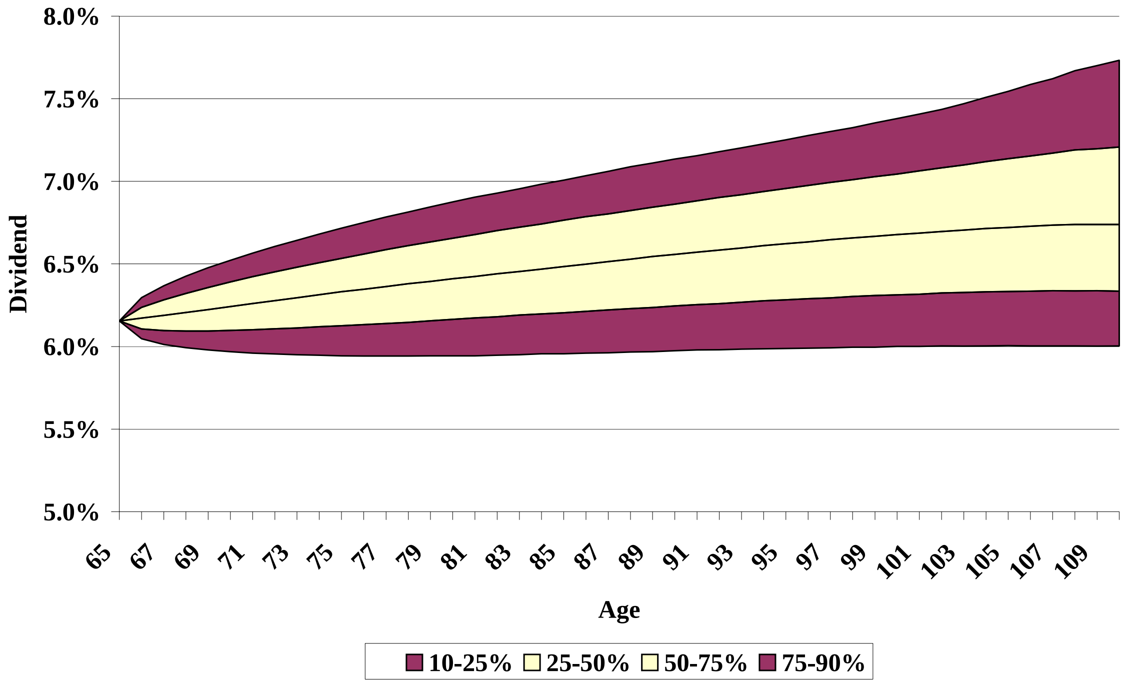

Figure 5: Case 3: Payment Rates on a Mortality-adjusted Annuity with $\alpha = 20$

Case 3: Allowance for risk and uncertainty. $\alpha = 20$ Figure 5 shows the results for the third case, with $\alpha = 20$, i.e. with a highly risk-averse population. The impact of precautionary saving is now plain. The initial payment is reduced to 6.15%. The median then drifts up over time to reach 6.68% at age 100 and 6.74% at age 110. At age 100 the lowest decile payment is 6.00% while the top decile payment is 7.38%. By age 110 the lowest decile remains at 6.00% while the top decile rises to 7.73%. Comparison of figures 4 and 5 indicates clearly the impact of precautionary saving resulting from risk aversion. There is a general upward drift in the pay-out pattern shown in figure 5 which is absent from figure 4. The mean shows upward drift similar to that of the medium.

Summary of Results We summarise the results of the three cases in table 1. This makes it clear that, even at the age of eighty-five, the effects of the different treatments of uncertainty and risk do not have large impacts. The fact that they emerge in extreme old age is, perhaps, not a very surprising implication of the underlying model used for two re–inforcing reasons. First, disturbances to mortality rates are logarithmic, so that given shocks have much larger absolute impact when the mortality rates themselves are large. Secondly, the error process adopted means that proportionate shocks to mortality, relative to what is known when the annuity is transacted are much larger extreme old age than in early old age. This compounds the first effect.

Age	Percentiles 10%	25%	50%	75%	90%
	Case 1: No uncertainty adjustment				
65	6.35	6.35	6.35	6.35	6.35
85	5.81	6.04	6.31	6.58	6.81
100	5.50	5.75	6.05	6.35	6.62
	Case 2: $\alpha = 0$				
65	6.33	6.33	6.33	6.33	6.33
85	5.82	6.01	6.33	6.60	6.84
100	5.71	5.99	6.32	6.67	7.00
	Case 3: $\alpha = 20$				
65	6.15	6.15	6.15	6.15	6.15
85	5.96	6.20	6.48	6.68	7.01
100	6.00	6.31	6.67	7.04	7.38

Table 1: Percentiles for Payment Rates on Mortality-adjusted Annuities by Age

An Extreme Case An alternative view of the implications of uncertainty and risk aversion can be gained by looking at extreme cases. We select the simulation which delivers the lower 1/2 percentile payment at age 80. We show in figure 6 the pay-outs which result in all three cases. The payment shows the expected pattern, being lower in the first part of the period and higher in the later part of the period when account is taken of risk aversion. The payment paid when $\alpha = 20$ settles at a figure about 0.2 percentage points higher than when $\alpha = 0$.

But it is also striking that, even when no account is taken of risk aversion, the payment falls eventually by only about 1/5th of the initial payment made when both risk and uncertainty are ignored. This indicates why high levels of risk aversion have apparently little effect. The impact of aggregate mortality risk on annuity payments is simply not that great, and therefore even highly risk-averse individuals are not prepared to do very much to protect themselves.

8 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 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.

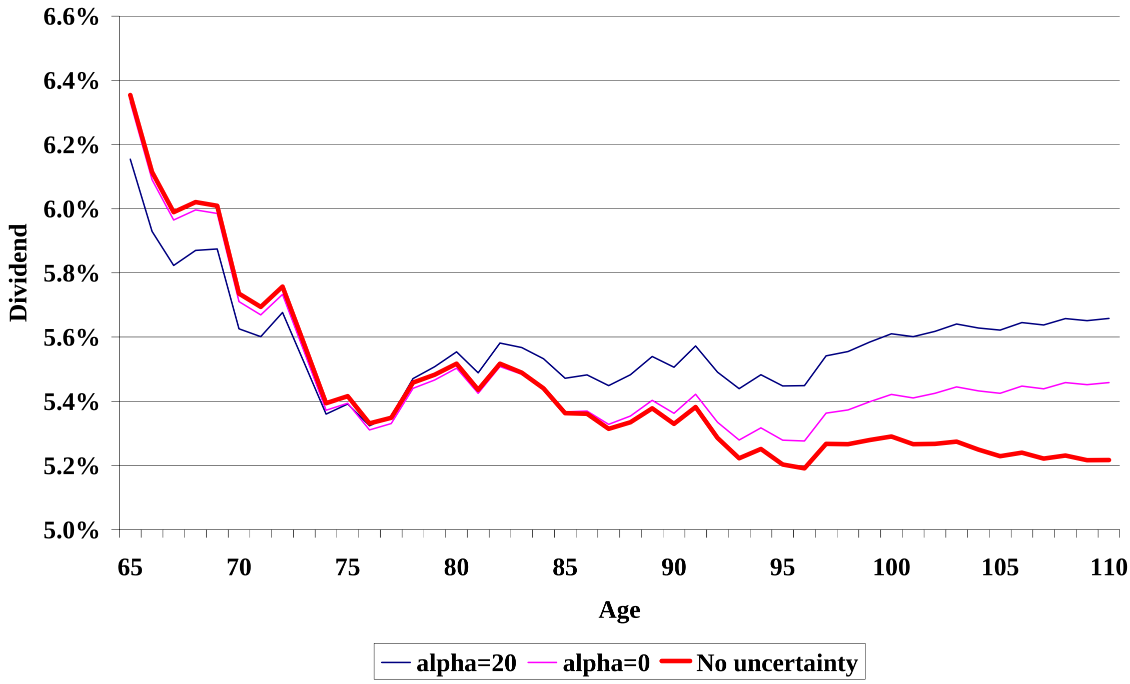

Figure 6: The Payment Profile of the Simulation which results in the lower 1/2 percentile Payment at Age Eighty

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.

Even, however, if one looks at an extreme case- the simulation which delivers the lower half-percentile payment at age eighty, the annuity rate is reduced by only about 1/5th when no account is taken of risk. This too suggests that risk-averse annuitants are not likely to regard carrying aggregate mortality risk to be a major problem.

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. But, as Mitchell (2001) shows, the presence of a stable source of income increases people's tolerance of the risk associated with other types of income. Thus people with modest privately-funded pensions have less need for protection from aggregate mortality risk than do those with larger retirement savings.

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.

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. 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.

Blake, D. & Burrows, W. (2001), 'Survivor Bonds : Helping to Hedge Mortality Risk', *Journal of Risk and Insurance* **68**, 339–348.

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.

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.

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

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

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.