



Network for Studies on Pensions, Aging and Retirement

Netspar DISCUSSION PAPERS

Joao Cocco and Francisco Gomes

**Longevity Risk, Retirement Savings, and
Financial Innovation**

DP 03/2011-061

Longevity Risk, Retirement Savings, and Financial Innovation

João F. Cocco¹

Francisco J. Gomes²

¹London Business School, Regent's Park, London NW1 4SA, CEPR, and Nestpar. Tel (020) 7000 8216.

Email jcocco@london.edu.

²London Business School, Regent's Park, London NW1 4SA, CEPR, and Netspar. Tel (020) 7000 8215.

Email fgomes@london.edu.

We would like to thank David Laibson, Bertrand Melenberg, Alex Michaelides, Olivia Mitchell, Thomas Post, Pedro Santa-Clara, an anonymous referee, and seminar participants at the SED meetings, the Conference on Financial Markets and Real Activity, London School of Economics, Stockholm School of Economics, Tilburg University, University of Aberdeen, and the Wharton School for comments. We also thank Andrew Cairns for making available his code for the estimation of the Renshaw-Haberman model.

This version: March 2011

Abstract

Over the last couple of decades there have been unprecedented increases in life expectancy which have raised important concerns for retirement savings. We solve a life-cycle model with longevity risk, which can be hedged through endogenous saving and retirement decisions. We investigate the benefits of financial assets designed to hedge the shocks to survival probabilities. When longevity risk is calibrated to match forward-looking projections those benefits are substantial. This lends support to the idea that such hedging should be pursued by defined-benefit pension plans on behalf of their beneficiaries. Finally, we draw implications for optimal security design.

1 Introduction

Over the last few decades there has been an unprecedented increase in life expectancy. For example, in 1970 a 65-year-old United States male individual had a life expectancy of 13 years.³ Roughly three and a half decades later, in 2007, a 65-year-old male had a life expectancy of 17.5 years. This represents an increase of 1.2 years per decade. To understand what such an increase implies in terms of the savings needed to finance a given stream of retirement consumption, consider a fairly-priced annuity that pays \$1 real per year, and assume that the real interest rate is 2 percent. The price of such annuity for a 65 year old male would have increased from \$10.5 in 1970, to \$13.5 by 2007. This is an increase of roughly 29 percent. Or in other words, to finance a given stream of real retirement consumption, a defined-benefit pension plan (or a 65-year-old male) would have needed 29 percent more wealth in 2007 than in 1970.

These large increases in life expectancy were, to a large extent, unexpected and as a result they have often been underestimated by pension plans, actuaries and insurers. This is hardly surprising given the historical evidence on life expectancy. From 1970 to 2007 the average increase in the life expectancy of a 65 year old male was 1.2 years/decade, but, from 1933 to 1970, the corresponding increase had only been 0.2 years/decade. This pattern of increases in life-expectancy has not been confined to the US. In the United Kingdom, a country for which a longer-time series of data on mortality is available, the average increase in the life expectancy

³The data in this paper on life-expectancy was obtained from the Human Mortality Database.

of a 65 year old male was 1.5 years/decade from 1970 to 2009, but only 0.1 years/decade from 1849 to 1970. These unprecedented longevity increases are to a large extent responsible for the underfunding of pay as you go state pensions,⁴ and of defined-benefit (DB) company and state sponsored pension plans, and to the recognition that either their benefits must be lowered or contribution rates increased.⁵

The response of governments has been to decrease the benefits of state pensions, and to give tax and other incentives for individuals to save privately, through defined contribution pension schemes. Likewise, many companies have closed company sponsored defined benefit plans to new members, while others have reduced their benefits and/or increased contribution rates. For instance, in their 2010 review of employer rates, the members of the benefits and program administration committee of Calpers note that the one assumption causing the biggest increase in employer contribution rates is the proposed change to more realistic post-retirement mortality assumptions.⁶

This also raises to question as to how should DB pension plans address the issue of longevity risk, or further changes in mortality rates going forward. It is true that the risk that plan members live longer may be reduced by the pension plan by the purchase of annuities at retirement age. However, for younger pension plan members there is considerable uncertainty

⁴The decrease in birth rates that has occurred over this period has also contributed to the underfunding.

⁵For individuals who are not covered by such defined-benefit schemes, and who have failed to anticipate the observed increases in life expectancy, a longer live span implies a lower average level of retirement consumption.

⁶<http://calpensions.com/2010/04/21/a-second-calpers-rate-hike-for-employers/>.

with respect to the level of aggregate life expectancy, and consequently the annuity prices, that they will face when they retire. How should pension plans address such risk? Should they try to hedge it to the extent that it is possible to do so? Naturally, the answer to these questions depends on the extent to which the pension plan beneficiaries might benefit from such insurance. This paper studies the extent to which individuals are affected by longevity risk, and the role that different instruments, including financial assets, play in hedging it. Thus, we focus on the demand for such assets.

We first document the existing empirical evidence on longevity, focusing on its historical evolution, on forward-looking estimates of mortality rates, and on the uncertainty surrounding these estimates. For this purpose we use current long-term projections made by the US Social Security Administration and by the UK Government Actuaries Department (GAD). We use this evidence to parameterize a life-cycle model of consumption and saving choices. The main distinctive feature of the model is that the survival probabilities are stochastic and evolve according to the specification proposed in Renshaw-Haberman (2006). This generalizes the one proposed in Lee-Carter (1992), which is the leading statistical model of mortality in the demographic literature, by allowing for cohort effects. For part of the analysis we focus on the simpler Lee-Carter formulation.

In our model, the individual receives a stochastic labor income each period, and decides how much to consume and save. She knows the current survival probabilities, but she does not know the future survival probabilities, since those are stochastic. Naturally the individual

forms an expectation of such probabilities when making her decisions. We allow for endogenous retirement so that, in addition to adjust her savings in response to changes in life expectancy, she can also revise her retirement decision.

Traditionally markets were incomplete in that agents did not have at their disposal the financial assets that would allow them to hedge longevity risk. We say traditionally since there have been recent attempts to address this market incompleteness. In December 2003, Swiss Re. issued a 400 million dollars three-year life catastrophe bond.⁷ This was a direct attempt by Swiss Re. to insure itself against a catastrophic mortality deterioration (e.g. a pandemic). More recently, there has been a growing interest in longevity swaps. These assets allow pension funds and other annuity providers to hedge the longevity risk to which they are exposed. In recent years several pension funds have actively started hedging their longevity exposure using financial products like these, provided by some of the major financial institutions. BusinessWeek reported a volume of 15 billion dollars in new life-settlement backed securities issued in the US in 2006, and this number was expected to double in 2007.⁸ According to Risk.net the volume of new issuances in the UK exceeded 7 billion pounds in 2009, with comparable numbers reported by The Financial Times.⁹

It is with this process in mind that we allow the agent in our model to invest in financial

⁷Blake and Burrows (2001) proposed the design of financial instruments for hedging longevity risk, and recommended that governments issue “survival bonds”, to allow the private sector to hedge this source of risk.

⁸BusinessWeek, July 30th, 2007.

⁹Financial Times, October 19th, 2009.

assets whose returns are correlated with the shocks to the survival probabilities, which we call longevity bonds. We study the portfolio allocation between these bonds and riskfree assets, namely how the demand for them changes over the life-cycle, and with individual characteristics. Therefore, our model allows us to identify, in a micro setting, who are the individuals who benefit most from longevity bonds, and those who benefit less, who might be the counterparty for such bonds.

We find that agents in our model respond to longevity improvements by increasing their savings, and in this way they are able to at least partially self insure against longevity shocks. Because longevity risk is realized slowly over the life-cycle, agents have time to react to the shocks. This of course requires that agents are well informed of the improvements in life-expectancy, and the implications of such improvements for the retirement savings needed. In addition, our model shows that retiring later is a typical response to improvements in life expectancy, even though such a decision carries a utility cost in terms of the foregone utility of (additional) leisure. Thus our model lends support to the argument that it is important to allow for flexible retirement arrangements as a mechanism for individuals to react to increases in life expectancy.

However, and importantly, our results show that even when agents are allowed to respond to shocks to life expectancy by saving more and by retiring later, longevity risk can have significant utility implications. More precisely, we find that individuals would benefit from being able to invest in longevity bonds, or financial assets whose returns are correlated with

longevity shocks. This is particularly the case when the extent of longevity risk in our model is calibrated to match the forward looking projections of the US Social Security Administration, and particularly those of the UK Government Actuaries Department.

Furthermore, the benefits from investing in longevity bonds are substantially higher when we take into account that the payouts of defined benefit pension plans are likely to decrease, and that this decrease is likely to be correlated with aggregate survival rates. This scenario is motivated by recent events, which suggest that as survival rates increase, retirement benefits are progressively decreased (or contributions increased). In this case, when longevity increases and households need more wealth to finance their retirement consumption, they are also more likely to receive a lower pension. Currently, there is not a liquid market for the longevity bonds that we model, but the partial-equilibrium utility gains associated with them do provide a meaningful metric of the extent to which individuals are affected by longevity risk under these different scenarios. This is metric that is also useful for DB pension plans deciding on whether to hedge longevity risk on behalf of members.

There is evidence of a link between income and life expectancy, with the larger mortality improvements occurring for individuals in the top half of the earnings distribution. We study the effects of such correlation, to find out that in our model individuals with higher income and life expectancy tend to accumulate more wealth. Finally, we use the model to study the optimal design of longevity bonds. Since investors face short-selling constraints, an increase in the volatility of the payoffs to the longevity bond allows them to achieve levels of hedging that

were previously unfeasible. However, excessive levels of volatility might deter young households from buying these asset, since they do not allow households to hedge labor income risk, which is their primary concern. This is an important trade-off to consider, when designing these securities.

There is a vast and rich literature on annuities, to which our paper is related. This literature has studied the welfare benefits that individuals derive from purchasing annuities, and why in spite of the large theoretical benefits of such purchases, in practice individuals do not seem to annuitize a significant part of their savings. Recent contributions to this literature include Mitchell, Poterba, Warshawasky, and Brown (1999), Brown, Davidoff, and Diamond (2005), Brown and Poterba (2006), Inkmann, Lopes and Michaelides (2008), Horneff, Maurer, Mitchell, and Stamos (2009), Chai, Horneff, Maurer, and Mitchell (2009), Yogo (2009) and Post (2011)).

In spite of being related to these papers, there are important differences between the financial asset that we study and annuities. First, in our model, longevity bonds are available for individuals to purchase even before retirement age, allowing individuals to hedge increases in life expectancy from early in life. In this respect alone this asset resembles more the purchase of forward annuities. Second, longevity bonds provide a hedge against aggregate life expectancy, not individual life expectancy. While this reduces the benefits for individuals wishing to hedge increases in their own life expectancy, it also means that, relative to annuities, adverse selection problems are reduced or eliminated. Therefore a market for longevity bonds could be more easily developed than a market for forward annuities, with individuals being able to in each

period buy and sell longevity bonds.

There is a growing literature that studies the optimal pricing of longevity bonds and related instruments (see, for example, Dahl (2004) and Carins, Blake and Dowd (2006)). In our paper we take bond prices as given and investigate their role in household portfolios, in the context of an empirically parameterized life-cycle model with labor income risk and endogenous retirement (see Farhi and Panageas (2007) for a model with endogenous retirement but without longevity bonds). We are interested in evaluating in a realistic setting the potential demand for longevity bonds by households saving for retirement, and consequently by pension providers acting on their behalf.¹⁰

The paper is organized as follows. In section 2 we use long term data for a cross section of countries to document the existing empirical evidence on longevity. In section 3 we setup and parameterize a life cycle model of the optimal consumption and saving choices of an individual who faces longevity risk. The results of the model are discussed in section 4. Section 5 explores the relation between life expectancy and income. Section 6 addresses security design implications. The final section concludes.

¹⁰Menoncin (2007) also introduces longevity bonds in an optimal savings and portfolio choice problem but in his model there is no labor income or retirement.

2 Empirical evidence on longevity

In this section we consider the existing empirical evidence on longevity. The data is from the Human Mortality Database, from the University of California at Berkeley. At present the database contains survival data for a collection of 28 countries, obtained using a uniform method for calculating such data. The database is limited to countries where death and census data are virtually complete, which means that the countries included are relatively developed.

We focus our analysis on period life expectancies. These life expectancies are calculated using the age-specific mortality rates for a given year, with no allowance for future changes in mortality rates. For example, period life expectancy at age 65 in 2006 would be calculated using the mortality rate for age 65 in 2006, for age 66 in 2006, for age 67 in 2006, and so on. Period life expectancies are a useful measure of mortality rates actually experienced over a given period and, for past years, provide an objective means of comparison of the trends in mortality over time. Official life tables are generally period life tables for these reasons.

Over the years there have been very significant increases in life expectancy at younger ages. For example, in 1960 the probability that a male US newborn would die before his first birthday was as high as 3 percent, whereas in 2000 that probability was only 0.8 percent. In England, and in 1850, the life expectancy for a male newborn was 42 years, but by 1960 the life expectancy for the same individual had increased to 69 years. We focus on life expectancy at ages 30 and 65 is due to the fact that we are interested on the relation between longevity risk and saving

for retirement. Furthermore, the increases in life expectancy that have occurred during the last few decades have been due to increases in life expectancy in old age. This is illustrated in Figure 1, which plots life expectancy for a male and female individuals for the United States over time, at birth, age 30, and at age 65.

Comparing the age-65 life expectancy of male and female individuals, we see that, as expected, it is higher for the latter group. More interesting is the fact that the age-65 life expectancy of females has increase more than that of males from 1970 to 2000, but the reverse is true in the latter part of the sample. Table 1 reports the average annual increases in life expectancy for a 65 year old individual for selected countries included in the database and for different time periods. It is important to note that the sample period available differs across countries. This table shows that there have been large improvements in life expectancy since 1970, and that they have not been confined to the US or England. Furthermore, except for Japan, there does not seem to be evidence that these increases are becoming smaller over time: in the US, and in the 1970s, the average annual increase was 0.13 years, whereas in the 1990s it was 0.1 years. In England, the corresponding values are 0.07 and 0.15. These improvements in life expectancy have been attributed to changes in lifestyle, smoking habits, diet, and improvements in health care, including the discovery of new drugs.

Figure 2 shows the conditional probability of death for a male US individual, for different years, and for ages 30 to 110. This figure shows, for each age, the probability that the individual will die before his next birthday. As it can be seen, the probability of death has decreased

substantially from 1970 to 2007, mainly after ages 65. This confirms the results in Figure 1, that the increases in life expectancy that have occurred over the past few decades have been due to decreases in mortality at old age.

The data reported in Table 1 is historical. However, we are mainly interested in forward looking estimates of future mortality improvements and in the uncertainty involved in such estimates. For this purpose we use data from two sources: the projections made by the US Social Security Administration and those made by the UK Government Actuary's Department (GAD). The latter were made in 2008, so that to facilitate comparison we use the projections provided in the 2008 OASDI Trustees Report. These projections were made taking into account not only the previously reported historical evidence, but also the opinion of experts in the field.

Panel A Figure 3 plots the projected future increase, in number of years, in the cohort life expectancy of a 65 year old male, relative to 2007.¹¹ Importantly, this figure plots for both the US and the UK an intermediate estimate, and a low and a high estimates. The UK GAD makes projections for every year whereas the US Social Security makes projections for every five years. We will use these to try to infer the future uncertainty in mortality improvements.

Comparing the US and the UK projections, we see that the US Social Security projects lower

¹¹Cohort life expectancies are calculated using age-specific mortality rates which allow for known or projected changes in mortality in future years. If mortality rates at a given age and above are projected to decrease in future years, the cohort life expectancy at that age will be greater than the period life expectancy at the same age.

expected increase in life expectancy, and lower dispersion in its estimates than the UK GAD. For instance, the UK GAD projects an intermediate (low/high) increase in life expectancy of a male between 2007 and 2045 of 3.8 years (0.6/8) whereas the corresponding numbers for the US are 2.5 years (1/4.5). There may be a variety of reasons for these differences, including the fact that the data refers to two different countries. However, Waldron (2005) carries out a comprehensive review of the long-term mortality projections available for the US, and finds that those made by the US Social Security Trustees are the most conservative in terms of assumed rate of mortality decline. Nevertheless, we will use both the UK and US projections reported in Figure 3 to parameterize the model.

3 The model

3.1 *Model setup*

3.1.1 *Survival probabilities*

We solve a life-cycle model of consumption and savings, in the spirit of Carroll (1997) and Gourinchas and Parker (2001), but in which survival probabilities are stochastic.¹² We let t

¹²In our model we consider shocks to aggregate survival probabilities that are exogenous to the agent. Yogo (2009) incorporates, in a model of the consumption and saving decisions of retirees, health investments through which agents can affect their own specific survival probabilities. His paper is concerned with matching portfolio

denote age, and assume that the individual lives for a maximum of T periods. Obviously, T can be made sufficiently large, to allow for increases in life expectancy in very old age. We use the Renshaw-Haberman (2006) model to describe the evolution of mortality rates. In this model mortality rates are given by:

$$\ln(m_{t,x}) = a_t + b_t \times k_x + c_t \times \gamma_{t-x}, \quad (1)$$

where $m_{t,x}$ is the death rate for age t in period x . The a_t coefficients describe the average shape of the $\ln(m_{t,x})$ surface over time. The b_t coefficients tell us which rates decline rapidly and which rates decline slowly in response to changes in the index k_x . The b_t are normalized to sum to one, so that they are a relative measure. The index k_x describes the general changes in mortality over time. If k_x falls then mortality rates decline, and if k_x rises then mortality worsens. When k_x is linear in time, mortality at each age changes at its own constant exponential rate. Finally, the γ_{t-x} term measures cohort-effects, i.e. the evolution of mortality rates for individuals born in different years. Naturally, this is a statistical model and therefore it does not specify the economic channels behind the cohort effects. One possibility, that we explore in section 5, is that these cohort effects are related to income.

This model reduces to the well known Lee-Carter (1992) model when the coefficients c_t are set equal to zero. Lee and Carter (1992) have shown that a random walk with drift describes

decisions and health expenditures.

the evolution of k_x over time well. That is:

$$k_x = \mu^k + k_{x-1} + \varepsilon_x^k, \quad (2)$$

where μ^k is the drift parameter and ε_x^k is normally distributed with mean zero and standard deviation σ^k . This model can be used to make stochastic mortality projections. The drift parameter μ^k captures the average annual change in k , and drives the forecasts of long-run change in mortality. A negative drift parameter indicates an improvement in mortality over time. For part of our analysis, we abstract from cohort effects, and focus on the simpler Lee-Carter model.

3.1.2 Preferences

Let p_t denote the probability that the individual is alive at age $t + 1$, conditional on being alive at age t . For a given individual age and time are perfectly co-linear, so that in order to simplify the exposition, and without loss of generality, we normalize x to be equal to t , and from now on we include only age indices. We allow the individual to choose when to retire. But while working she incurs a cost in leisure time, which we denote L_t . Furthermore, we assume that the individual's preferences are described by the time-separable power utility function:

$$E_1 \sum_{t=1}^T \delta^{t-1} \left(\prod_{j=0}^{t-2} p_j \right) \left\{ p_{t-1} \frac{(L_t^\gamma C_t)^{1-\theta}}{1-\theta} + b(1-p_{t-1}) B(D_t, p_t) \right\}, \quad (3)$$

where δ is the discount factor, C_t is the level of age/date t consumption, θ is the coefficient

of relative risk aversion, γ measures the preference for leisure, and D_t is the amount of wealth the individual bequeaths to his descendants at death. The parameter b controls the intensity of the bequest motive, and $B(D_t, p_t)$ describes the utility from leaving a bequest. We will consider different alternatives for this function. To simplify and reduce the number of choice variables, we let L_t be equal to two-thirds while the individual is working, and equal to one during retirement.

3.1.3 Labor income

During working life age- t labor income, Y_t , is exogenously given by:

$$\log(Y_t) = f(t, Z_t) + v_t + \eta_t \quad \text{for } t \leq t_R, \quad (4)$$

where $f(t, Z_t)$ is a deterministic function of age and of a vector of other individual characteristics, Z_t , and two random components (as standard in the life-cycle literature, e.g. Viceira (2001) or Gourinchas and Parker (2002)) η_t is an idiosyncratic temporary shock distributed as $N(0, \sigma_\eta^2)$, and v_t is a permanent income shock, with $v_t = v_{t-1} + u_t$, where u_t is distributed as $N(0, \sigma_u^2)$ and is uncorrelated with η_t . We let shocks to permanent income u_t be correlated with shocks to longevity ε_x^k (since for a given individual time and age are co-linear $t \equiv x$), and denote by ρ the corresponding correlation coefficient. Thus before retirement, log income is the sum of a deterministic component that can be calibrated to capture the hump shape of earnings over the life cycle, and two random components, one transitory and one persistent.

Retiring later in life may be an additional natural mechanism to insure against increases in life expectancy. Therefore we allow the individual to choose when to retire, and we let t_R denote this endogenously chosen retirement age.¹³ After this age income is modeled as a constant fraction λ of permanent labor income in the last working-year:

$$\log(Y_t) = \log(\lambda) + f(K, Z_K) + v_K \text{ for } t > t_R . \quad (5)$$

The parameter λ measures the extent to which the individual has defined benefit pensions, which implicitly provide insurance against longevity risk. To simplify the solution to the problem, we assume that retirement is a once and for all decision, that the individual may not retire before a minimum retirement age t_R^{min} , nor after a maximum retirement age t_R^{max} .

3.1.4 *Financial assets*

We assume that there are two financial assets in which the individual can invest. The first is a riskless asset which has interest rate R . The second is a financial asset, which we call longevity bond, whose returns are perfectly negatively correlated with innovations to mortality, thus providing the investor with a perfect hedge against this risk. More precisely, we assume that the return on longevity bonds (R_t^L) is given by:

$$R_t^L = \mu^L - \frac{\sigma^L}{\sigma^k} \varepsilon_t^k, \quad (6)$$

¹³Farhi and Panageas (2007) solve for the optimal retirement age in the context of a life-cycle model of consumption and portfolio choice.

where μ^L and σ^L are the mean and standard deviation of longevity bond returns respectively, and ε_t^k is the age t (and time x since, in this case, $t \equiv x$) shock to the mortality rates. Recall that a negative value for ε_t^k implies an improvement in life expectancy, so that when this happens bond returns are high.

There are important differences between longevity bonds, as we model them here, and annuities, an asset which has been studied extensively in the literature. First, longevity bonds provide investors with an hedge against aggregate mortality shocks. Second, longevity bonds can be purchased by investors in our model in each period, even when young, allowing them to obtain insurance against the additional savings that future increases in life expectancy will require.¹⁴

3.1.5 *The optimization problem*

In each period the timing of the events is as follows. The individual starts the period with wealth W_t . Then labor income and the shock to survival probabilities are realized. Following Deaton (1991) we denote *cash-on-hand* in period t by $X_t = W_t + Y_t$. We will also refer to X_t as wealth: it is understood that this includes labor income earned in period t . Then the individual must decide how much to consume, C_t , which fraction of savings to invest in the riskless asset and in longevity bonds, and for periods in between the minimum and maximum retirement age

¹⁴In addition to annuities, it would be interesting to expand the set of assets considered to include equities and the more traditional bonds, as in Campbell, Chan and Viceira (2003) and Munk and Sorensen (2010).

whether to retire, if she has not done so before. The wealth in the next period is then given by the budget constraint:

$$W_{t+1} = (1 + R_{t+1}^p)(W_t + Y_t - C_t), \quad (7)$$

where

$$R_{t+1}^p = \alpha_t R_{t+1}^L + (1 - \alpha_t)R, \quad (8)$$

and α_t is the share of wealth invested in longevity bonds at time t . The problem the investor faces is to maximize utility subject to the constraints. The control variables are consumption/savings, the proportion of savings invested in the longevity bond, and whether or not to retire. The state variables are age, cash-on-hand, the current survival probabilities, and a dummy variable which takes the value of one if the investor is currently working, and zero otherwise. In our setup the value function is homogeneous with respect to permanent labor income, which therefore is not a state variable.

3.2 *Calibration*

3.2.1 *Survival probabilities*

Undoubtedly the calibration of the parameters for the mortality process is likely to be the most controversial. This is in itself a sign that there is a great deal of uncertainty with respect to what one can reasonably expect for future increases in life expectancy. In this section we abstract from cohort effects and focus on the simpler Lee-Carter model. The reason why we

have decided to do so is that when we have estimated the RH model we have found that there is some colinearity between the age coefficients, period (k_x) and cohort (γ_{t-x}) effects, which led to some of the coefficients being imprecisely estimated, and the model results being somewhat dependent on the exact cohort that we solve the model for. Nonetheless, in section 5 we show results for the more general RH model.

In order to parameterize the stochastic process for survival probabilities we follow two alternative approaches. First, we estimate the parameters of the Lee-Carter model using *historical* data. Second, we try to determine which are the parameters of such model that match the projected increases in life expectancy shown in Figure 3, made by the US Social Security and the UK Government Actuaries Department. The latter projections are forward looking measures that combine historical data with other information, namely expectations of future improvements in mortality.

For the estimation of the Lee-Carter model using historical data, we use US data from 1933 to 2007, which is the data period available in the Human Mortality Database, and estimate:

$$\ln(m_{t,x}) = a_t + b_t \times k_x + \zeta_{t,x}, \tag{9}$$

where $\zeta_{t,x}$ is an error term with mean zero and variance σ_ζ^2 , reflecting particular age-specific historical influences not captured by the model. Naturally, in the data the model will not hold exactly, so for estimation purposes we need to add a residual. For most of the paper we will ignore it and consider the model as holding exactly, but in section 4.2 we report some model

results when we allow for uncertainty in the estimated parameters. This model is undetermined: k_x is determined only up to a linear transformation, b_t is determined only up to a multiplicative constant, and a_t is determined only up to an additive constant. Following Lee and Carter (1992) we normalize the b_t to sum to unity and the k_x to sum to zero, which implies that the a_x are the simple averages over time of the $\ln(m_{t,x})$. This model cannot be fit by ordinary regression methods, because there are no given regressors. On the right side of the equation there are only parameters to be estimated and the unknown index k_x . We apply the singular value decomposition method to the logarithms of the mortality rates after the averages over time of the log age-specific rates have been subtracted to find a least squares solution.

Figure 4 shows the actual and estimated mortality rates for two different years, namely 1950 and 2007. From this figure we see that the model fits the data relatively well, although the fit is worse in the early year and at advanced ages. We use the time series data of k_x to estimate the parameters of the random walk. For males, the estimated drift parameter μ^k is -0.65 and the standard deviation of the shocks σ^k is 1.86 (both of these parameters are reported in Table 2.1). For females, the corresponding historical estimates are -0.90 and 1.82 , respectively, so that their experience was characterized by larger mean improvements but with slightly lower volatility.

In order to understand what such estimated parameters imply in terms of future improvements in life expectancy, we use them to make projections. More precisely, we ask the following question: consider an individual who is 30 years old in 2007 (the starting age in our model

and the latest year for which we have data from HMD database respectively). How does life expectancy at age 67 in 2007 compare to life expectancy at age 67 when the individual reaches such age (i.e. in year 2045)? In other words, we ask which is the increase in the life expectancy of a 67 year old individual that the model forecasts to take place over the next 37 years?¹⁵

Obviously, such increase will be stochastic as it will depend on the realization of the shocks to survival probabilities that will take place over the next 37 years. Therefore, in Table 2.2 we report increases/decreases in life expectancy for several percentiles of the distribution of the shocks to life expectancy (25, 50, and 75). The first row of Table 2.2 shows that, for the values of μ^k and σ^k estimated using historical data, the median increase in the life-expectancy of a 65 year old over such a period is 2.3 years. The 25th and 75th percentiles are 1.6 and 2.9 years, respectively.

The calculations in the first row of Table 2.2 are based on a statistical model that extrapolates for the future based on the history of mortality improvements that has taken place between 1933 and 2007, without conditioning on any other information. We carry out a second calibration exercise in which we choose parameters for the Lee-Carter model such that when we simulate the model we are able to generate increases in life expectancy that roughly match the forward looking projections of the US Social Security and of the UK GAD shown in Figure 3. In panels A2 and B2 of Table 2.1, and based on the same data that we have used in Figure

¹⁵The reason why we focus on life expectancy at age 67 and not 65 is that the US social security projections are only provided for every five years.

3, we report the US Social Security and the UK GAD projected increases in the life expectancy of a 67 year old over a 37 year period (from 2007 to 2045). We report the projected increase for the low, principal, and high variants/cost.

These intermediate, high and low cost projections are not necessarily carried out in the context of a model, but they agreed upon by the government actuaries. For the UK, they are calculated by assuming annual rates of mortality improvement of 1, 0.5 and zero percent at all ages for the high, principal and low life-expectancy variants, respectively. Therefore, it is not possible to assign probabilities to these different variants. The GAD reports that “these (high and low variants) are intended as plausible alternative assumptions and do not represent lower and upper limits.” (GAD report no. 8, page 28) It is not clear what “plausible alternative assumptions” means exactly, but we think that it is reasonable to compare the high and low variant projections to the 25th and to the 75th percentiles of the distribution of improvements in life expectancy that is generated by the model.

Comparing the mortality improvements projected using the estimated historical parameters of the Lee-Carter model (shown in Table 2.2), to the projections of the US Social Security and the UK GAD (shown in Panels A2 and B2) we see that the former lead to lower dispersion in the forecasts, particularly so when compared to the projections of the UK GAD.¹⁶ To motivate this

¹⁶One could also question the extent to which the projections made by the GAD for the UK, are comparable to the estimates for the US. However, Deaton and Paxson (2001) have shown that the UK and the US have similar histories of mortality improvements.

dispersion the GAD reports that “it could be argued that uncertainty over long-term mortality levels is higher than ever given current research in areas such as mapping the human genome and gene therapy.” (GAD report no. 8, page 25)

We then calibrated the drift and volatility parameters for the Lee-Carter model so as to try to match the US Social Security and the UK GAD projections. The second and third rows of Panel A1, Table 2.2 (Panel A2, for females) shows the values that we have obtained (Figure 5 plots the implied distributions). In the third column of Table 2.1 we report values for the drift and volatility parameters that generate mortality improvements that roughly match those projected by the US Social Security. These parameters involve a smaller drift parameter (for females) and a higher volatility than those estimated from historical data. The increase in volatility are much larger if we wish to roughly match the UK GAD projections instead (Table 2.1).

To somehow try to reconcile the historical estimates with the US Social Security projections we have also made mortality forecasts using the historical estimates but taking into account that there is uncertainty in the estimate of μ^k (by computing the posterior distribution based on a diffuse prior). While the forecasts based on the exact point estimates shown in the first row of 2.2 imply a narrow range of future improvements in life expectancy (between 1.6 and 2.9 years over a thirty seven year period), when we allow for parameter uncertainty in μ^k this range increases considerably (to between 0.5 and 4.3 years). The latter are comparable to the ones implied by US Social Security projections (between 1.0 and 4.5 years, for males, reported

in the first row of Panel A2, Table 2.2). Therefore, as a baseline case we use the parameters that better match the US Social Security projections.

3.2.2 *Preference parameters*

The initial age in our model is 30, and the individual lives up to a maximum of 110 years of age. That is T is equal to 110. The minimum retirement age, t_R^{min} , is set equal to 65, which is the typical retirement age. We assume that the maximum retirement age, t_R^{max} , is 70, so that the individual must retire at this age, if she has not done so before. Obviously, our model could accommodate other values for the minimum and maximum retirement ages, but we think that these are reasonable values. We assume a discount factor, δ , equal to 0.98, and a coefficient of relative risk aversion, θ , equal to three. We follow Gomes, Kotlikoff and Viceira (2008) and set γ equal to 0.9. In the baseline model we assume that there is no bequest motive. These values are reported in Table 3.

3.2.3 *Labor income and asset parameters*

To calibrate the labor income process, we use the parameters estimated by Cocco, Gomes and Maenhout (2005) for individuals with a high school degree, but with the dollar numbers converted to 2007 US dollars using the consumer price index. In the baseline case we assume zero correlation between labor income shocks and shocks to longevity, but in section 5 we

relax this assumption to explore the links between the two. One important possibility that we also consider, is that the retirement replacement ratio is correlated with improvements in life expectancy. This is motivated by recent events: the large improvements in aggregate life expectancy that have occurred over the last decades have led governments to reduce the benefits of pay as you go state pensions. Therefore we will consider a parameterization in which the replacement ratio is reduced when there is an improvement in life expectancy, such that the expected present value of the retirement benefits that the individual receives is unchanged.

We assume that the real interest rate is equal to 1.5 percent. The calibration of the returns on the longevity bonds is harder, since a market for these bonds as we have defined them, with returns perfectly correlated with an aggregate mortality index, does not exist. There have been attempts to issue assets with longevity-linked payoffs. For example, the bonds issued by Swiss Re. paid a quarterly fixed coupon equal to 3-month US dollar LIBOR plus 135 basis points. The principal was then repaid in full if the mortality index did not exceed a given threshold, but the payment decreased linearly with the index if that threshold was reached.

Friedberg and Webb (2005) compute hypothetical returns for longevity bonds, based on a specific issuance made by the European Investment Bank. They find these returns to be negatively correlated with aggregate consumption and, in a CCAPM framework with relative risk aversion of 10, the implied risk premium would then be -0.02% . They also estimate a return volatility of 3%. We use these values to parameterize our model so that we set the risk premium on longevity bonds to zero (that is $\mu^L = 1.5\%$). However, we also recognize that

these are fairly hypothetical calculations, and in section 6 we show results for alternative values for the first two moments of longevity bond returns (based on results from Cairns, Blake and Dowd (2006)). In that section we also discuss who the potential counterparties in such asset could be. On one hand, demand for longevity bonds is likely to come from pension plans. On the other hand, as illustrated by the Swiss Re. case, supply may arise from life insurers.¹⁷

3.3 *Solution technique*

The model was solved using backward induction. In the last period the policy functions are trivial (the agent consumes all available wealth) and the value function corresponds to the indirect utility function. We can use this value function to compute the policy rules for the previous period and given these, obtain the corresponding value function. This procedure is then iterated backwards.

To avoid numerical convergence problems and in particular the danger of choosing local optima we optimized over the space of the decision variables using standard grid search. The sets of admissible values for the decision variables were discretized using equally spaced grids. The state-space was also discretized and, following Tauchen and Hussey (1991), we approximated the density function for labor income shocks using Gaussian quadrature methods, to perform the necessary numerical integration.

¹⁷The answer to this question requires a general equilibrium model, and the contribution of our paper is to study the potential demand for longevity bonds in a realistic life-cycle setting.

In order to evaluate the value function corresponding to values of cash-on-hand that do not lie in the chosen grid we used a cubic spline interpolation in the log of the state variable. This interpolation has the advantage of being continuously differentiable and having a non-zero third derivative, thus preserving the prudence feature of the utility function. The support for labor income realizations is bounded away from zero due to the quadrature approximation. Given this and the non-negativity constraint on savings, the lower bound on the grid for cash-on-hand is also strictly positive and hence the value function at each grid point is also bounded below. This fact makes the spline interpolation work well given a sufficiently fine discretization of the state-space.

4 Results

4.1 *Life-cycle profiles*

In Figure 6, Panel A we plot the average simulated income, wealth and consumption profiles in the model. The patterns are very similar to the ones obtained in standard life-cycle models of consumption and savings decisions (e.g. Gourinchas and Parker (2002)). Early in life households are liquidity constrained and consumption tracks income very closely, with a small level of savings being accumulated to use as an insurance-cushion against negative labor income shocks. As the agent gets older the level of labor income increases, the profile becomes less steep, and

the agent starts saving for retirement.

As the agent approaches retirement age the consumption profile becomes less steep, and its level exceeds current income. The household is now partly consuming out of her previously-accumulated wealth: net savings become negative and wealth begins to fall. Since in our model the agent also derives utility from leisure, we are able to replicate the empirically observed “discrete” drop in consumption at retirement. In order to smooth intertemporal marginal utility, and to compensate for the desutility of working, the household increases consumption during working life. After retirement, and as mortality risk increases, the consumption path decreases at a relatively fast pace until it reaches the level of retirement benefits. Finally, Panel A of Figure 6 also plots the proportion of individuals working. This proportion declines from age 65 onwards, and almost all individuals are retired by age 70, which is the latest retirement age in our model.

In order to better understand the effects of longevity shocks on individual choices we compare the consumption/saving and retirement decisions of individuals under different realizations for the shocks to life expectancy. More precisely, we compare the choices of individuals who throughout working life face different magnitude of improvements in life-expectancy. It is important to note that this comparison across individuals should be seen as different possible future realizations of the aggregate shocks to life expectancy, and not as different concurrent realizations, since the longevity shocks that we model are aggregate.

More precisely in Figure 6, Panel B we report simulated age 65 cash-on-hand as a function of the different realizations of the shocks to life expectancy. From this figure we see that those individuals who throughout working life have faced positive shocks to life expectancy, so that age 65 life expectancy is higher, also have (on average) accumulated higher levels of savings. Naturally, one optimal response of agents to increases in life-expectancy is to save more. In this way, individuals are able to at least partially self-insure themselves against the fact that they expect to live longer, and need to finance more retirement consumption out of accumulated savings. Because longevity risk is realized slowly over the life-cycle, agents have time to react to the shocks. Thus the first implication that we can draw from our model, is that it is important that individuals, throughout their lives, are well informed of improvements in life expectancy. This result is suggestive of the importance of household financial literacy, a point emphasized by Lusardi and Mitchell (2006).

In addition, individuals in our model can also decide at which age to retire. Individuals with low life expectancy tend to retire at age 65, while those individuals with longer life expectancy tend to keep on working. The average conditional life-expectancy among those retired at age 65 is 16.5 years while for those who are still working it is 23.7 years. Therefore, individuals in our model also react to improvements in life expectancy by retiring later. Thus our model lends support to the argument that it is important to allow for flexible retirement arrangements as a mechanism for individuals to deal with increases in life expectancy.

Figure 7 plots the share of wealth invested in the longevity bond over the life-cycle. Early

in life the demand for this asset is crowded-out by labor income risk. Young households have not yet accumulated significant wealth and therefore are significantly exposed to labor income shocks. The riskless asset is better suited to hedge this risk and therefore these households prefer to invest a significant fraction of their wealth in riskless bonds. However, as they accumulate more wealth over time and start saving for retirement, the allocation to longevity bonds quickly increases, and approaches 100%. During retirement the portfolio allocation to longevity bonds remains high, but it soon becomes undetermined since household wealth converges to zero.¹⁸

4.2 *Utility gains*

We compute the utility gains from being able to invest in longevity bonds. This provides us with a measure of how much individuals are affected by longevity risk. We calculate these utility gains under the form of consumption equivalent variations (similarly to Balduzzi and Lynch, 1999). For each of the scenarios (with and without access to longevity bonds), we compute the constant consumption stream that makes the individual as well-off in expected utility terms as the one that can be financed by his decisions. The utility gain of having access to longevity bonds is the dollar difference in the present value of the (constant) consumption stream in the case the individual has no access to longevity bonds and the present value of the (constant) consumption stream in the case the individual has access to longevity bonds. We

¹⁸In the appendix we include a three period model that provides intuition on the asset allocation effects at work in the model.

use the riskfree rate to calculate the present value of this difference.¹⁹

It is important to point out that the utility measure that we calculate is an ex-ante utility measure, in the sense that it is a utility measure calculated as of age 30, for a given longevity risk. More precisely we compare the utility gains between two investors, with and without access to longevity bonds, that face the same stochastic process for longevity risk and the same life-expectancy (as well as the same other parameters). Thus we are not comparing individuals with different preference parameters (or who face different stochastic processes for longevity risk).

The benefits of investing in longevity bonds occur late in life, i.e. after retirement. The present value of such benefits, when discounted to age 30, does not provide us with a good idea of the relative importance of longevity bonds for financing retirement consumption. Therefore, in addition to reporting the present value of the utility gains as of age 30, we also calculate the utility gains, as of age 65, as a fraction of accumulated retirement wealth. This number is obtained simply by multiplying the age 30 utility gains by $(1 + R_f)^{35}$, i.e. capitalizing the certainty-equivalent gains to age 65, and dividing it by age 65 average accumulated retirement wealth under the scenario in which the agent does not have access to longevity bonds. Therefore, this second measure of utility gains is simply a re-scaling of the former.

The results are shown in Table 4. For males, and for the parameters that match the forward-

¹⁹Alternatively, we could have calculated the utility losses in terms of the percentage gain/loss in the consumption stream.

looking US SS projections (Lower volatility and lower mean improvement scenario), the utility gains of investing in longevity bonds are 343 US dollars measured at age 30, or 2 percent of age 65 retirement wealth. These gains are substantially larger when we parameterize the longevity process to match the less conservative UK GAD projections (Higher volatility and higher mean improvement scenario). For such parameters the utility gains are 947 dollars as of age 30, or as high as 5.6 percent of age 65 retirement wealth. However, the gains of investing in longevity bonds are considerably smaller when longevity risk is calibrated to match the historical data.

Therefore, the conclusion to be drawn from these numbers, is that if longevity risk going forward is similar to what can be inferred statistically from historical data and assuming that the Lee-Carter model is appropriate, the benefits from investing in longevity bonds are small. However, they become economically significant if longevity risk going forward is higher, as implied by the forward looking estimates of the US Social Security and especially the UK GAD projections. This is the case even after allowing for an endogenous retirement decision and optimal saving and consumption decisions as a response to longevity shocks.

One way of (at least partially) reconciling the results based on historical data with those based on parameters that better match the US Social Security projections is by recognizing that the historical estimates are subject to parameter uncertainty. We have extended our model to do so. More precisely, we have considered an extended version of the investor's optimization problem in which we have incorporated uncertainty about the value of μ^k (assuming a diffuse prior, and ignoring learning). We have solved this more general model for the estimated value

for volatility (and the other model parameters) from historical data alone, but with this added element of parameter uncertainty about the mean improvements in mortality. Since this is computationally demanding we have only replicated our baseline experiments, to assess the welfare gains from investing in longevity bonds in such a context. We have found that these gains are as high as 50% of the gains that we obtain when instead we calibrate the model to match the US Social Security forward-looking projections. Naturally, we would obtain a much higher value if we allowed for uncertainty in the other estimated parameters.

4.3 *Alternative retirement benefits*

In recent years there has been a trend away from defined benefit pensions, and towards pensions that are defined contribution in nature. Faced with the severe projected underfunding of pay as you go state pensions systems, many governments have reduced the level of benefits of such schemes, or are planning to do so. In addition, many companies have closed their defined benefit schemes to new employees, or have reduced the benefits that they are entitled to receive precisely as response to the underfunding of schemes. This is likely to exacerbate the longevity risk that individuals face. Not only are the benefits on average lower, but they are also reduced precisely when individuals expect to live longer.

We use our model to investigate these scenarios. More precisely, we carry out two different exercises. In the first we decrease the replacement ratio from the baseline value of 0.68 to 0.40

(based on the recent numbers from Caggeti and De Nardi (2006)), and investigate the effects of such lower replacement ratio. The results are shown in the last two columns of Table 4. The age 30 utility gains associated with the longevity bonds are considerably higher, and equal to 1025 and 2770 dollars for the US and the UK, respectively. As a fraction of the age 65 accumulated wealth these welfare gains are slightly smaller than the baseline scenario. This is of course due to the fact that the agents accumulate more wealth when the replacement ratio is lower.

In the second, and probably more realistic scenario, we allow for a stochastic replacement ratio, which is assumed to be negatively correlated with longevity shocks. In this case, when there are improvements in mortality rates the level of retirement benefits is decreased, so that the present value of all future retirement benefits is kept constant. This is motivated by recent events: the providers of defined benefit pension plans have in recent years reduced the benefits that are to be paid out as a response to the large increases in life expectancy that have occurred over the last decades.

In this scenario, the benefits for males (females) from having access to longevity bonds are an order of magnitude higher than before, and equal to 5.15 (5.35) percent and 9.62 (7.68) percent of age 65 retirement wealth, for the US SS and the UK GAD projections scenarios, respectively, and for a replacement ratio equal to 0.68. These are very significant gains and suggest that the hedging of longevity risk should be a major consideration for DB pension plans managing their assets on behalf of pension plan beneficiaries.

4.4 *Bequest motive*

In this section we consider the effects of a bequest motive on the benefits of longevity bonds. We consider two alternative specifications for this motive.²⁰ In the first, the individual derives utility from the wealth that he bequeaths, according to a CRRA utility. More precisely, the function $B(D_t, p_t)$ is assumed to be:

$$B(D_t, p_t) = \frac{D_t^{1-\theta}}{1-\theta}, \quad (10)$$

where D_t denotes the dollar amount bequeathed. In this specification the utility derived from leaving a bequest at time/age t does not depend on aggregate survival probabilities at this age. We call this a “joy of giving” bequest motive.

The fourth and fifth columns of Table 5 show the welfare gains of longevity bonds when the parameter that measures the intensity of the bequest motive (b) is equal to one, and for a replacement ratio of 40%. For comparison, the previous two columns of the same table report the results for no bequest motive scenario (previously shown in Table 4). Comparing the two, we see that the dollar value of the utility gains from investing in longevity bonds are now smaller. The intuition is simple: when there is a joy of giving bequest motive, the individual saves for his own retirement and to leave a bequest. An improvement in life expectancy means

²⁰We are grateful to an anonymous referee for encouraging us to investigate the effects of alternative bequest motive specifications.

that more saving is required for retirement but, on the other hand, less saving is required to finance the bequest motive (it takes place further into the future and so it is discounted more).

However, this result depends crucially on the joy of giving specification for the bequest motive. To see this consider an alternative altruistic bequest specification in which the individual cares directly about the utility of the descendent. More precisely, assume that at the time of the death of the parent, the child is 30 years old, and that the wealth bequeathed is converted into an annuity using the state contingent survival probabilities. This annuity is added to the child's per period income, whose lifetime utility enters the utility function of the parent. That is: the function $B(D_t, p_t)$ is simply the expected lifetime utility of the child, which depends on the state-contingent survival probabilities that prevail at the time of the death of the parent.²¹

In this alternative specification a shock to mortality rates which increases the life expectancy of the parent will also increase the life expectancy of the descendent, so that a given amount of wealth bequeathed corresponds to a smaller annuity value. Thus, when life expectancy is unexpectedly high the individual has to save more both for his/her own retirement and also for his/her children. The welfare gains of longevity bonds in this context are shown in the last two columns of Table 5. Their dollar values are now larger than under the no bequest motive specification. With an altruistic bequest motive individuals save considerably more than when such motive is not present, which explains why as a percentage of retirement wealth the benefits

²¹To simplify we calculate this lifetime utility assuming that the child consumes in each period his/her own income plus the annuity value that he/she has received from his/her parent.

of longevity bonds are not always higher than in the no bequest motive scenario.

5 Income and longevity

The recent increases in life-expectancy have been concentrated in the upper half of the income distribution. In a recent study Waldron (2007) presents an analysis of trends in mortality differentials and life expectancy by average relative earnings for male Social Security-covered workers aged 60 or older. She finds that the top half of the average relative earnings distribution has experienced much faster mortality improvement than the bottom half. As Waldron acknowledges, this and other studies that try to estimate the relationship between mortality and income are subject to data limitations and more importantly, endogeneity problems. The concern is that individuals who experience adverse health shocks (which increase their mortality probability) are more likely to drop out of the labor market, thus simultaneously suffering a drop in income. Thus, the problem is one of reverse causality.

In order to somehow try to address it, Waldron (2007) and Cristia (2007) use measures of lifetime average earnings to classify individuals (instead of single year), and compute these average measures excluding earnings observations in the years immediately preceding the year when mortality is ascertained. For example, Waldron (2007) uses average earnings between ages 45 and 55, and then studies mortality rates from age 60 onwards. Both of these papers find that the relationship between mortality and lifetime earnings is very strong, suggesting

that it is indeed present in the data.

We use our model to investigate the effects of longevity risk, when there is a relationship between income and mortality. We model this relationship in two different ways. In the first, we assume all individuals are ex-ante identical, but permanent income shocks are positively correlated with longevity shocks, so that those individuals who have higher income tend to have higher longevity. In the second we assume that individuals are ex-ante heterogeneous, and that those who have higher expected levels of lifetime income also have higher life expectancy.

5.1 *Correlated income and longevity shocks*

Within the Lee-Carter model we allow the shocks to mortality improvements to be correlated with the shocks to permanent income. Recall that in the baseline case we had set this correlation to zero. Here we calibrate it so that the simulated mortality rates across income groups closely matches the differences in mortality between the income groups reported by Waldron (2007) and Cristia (2007). For the former we match the ratios of age 65 to 69 mortality rates for individuals in the bottom and top half of the earnings distribution. The correlation coefficient that provides such match is 0.4. For the latter, we match the ratio of age 65 to 75 mortality rates for individuals in second and fourth income quintiles of the income distribution. This implies a correlation coefficient of 0.3. Therefore we have decided to use an intermediate value, equal to 0.35, to calibrate the model.

Table 6 reports the benefits of longevity bonds with correlated income and longevity shocks. From an ex-ante point of view, i.e. from the perspective of an individual that is age 30, the benefits of longevity bonds are now slightly smaller. For example, for a male facing lower longevity risk and a fixed replacement ratio of 0.68 the percentage gains of longevity bonds are 3.48 percent of retirement wealth for the case of positive correlation, compared to 5.56 percent in the absence of correlation (Table 4). Intuitively, when labor income shocks and mortality improvements are positively correlated, labor income is a (partial) hedge against longevity shocks. Therefore, from an ex-ante point of view, the benefits of longevity bonds are smaller.

5.2 *Income levels and longevity*

In the previous section all individuals are equal from an ex-ante (age 30) point of view. But it is important to recognize that individuals may be ex-ante heterogenous with some having higher expected lifetime earnings and, in the presence of correlation between income and mortality rates, they also have higher life expectancy. The RH model is well suited to investigate this issue, since we can interpret the cohort effects as arising from differences in income or lifetime earnings. A further advantage of this approach is that, when estimating the model for the different cohorts, we avoid endogeneity problems that may arise as a result of, for example, individuals with poor health dropping out of the labor market. Therefore, we follow the Deaton and Paxson (1999) approach of identifying exogenous (to the individual) sources of income

variation (such as cohort effects).²² The disadvantages are that the cohort effects are also likely to capture other factors in addition to income, and that there may be more variation in individual-level income and mortality than that estimated from cohort data.

We use the mortality data for the US from 1933 to 2007 to estimate the RH model. One of the difficulties that we encountered is that there is a large degree of colinearity between the age coefficients, period and cohort effects, which leads to some of the coefficients being only imprecisely estimated. This has also been documented by Cairns et. al (2009). Nonetheless, we have selected two (of the more precisely estimated) cohorts who were born two decades apart, and have forecasted their survival rates. Figure 8 plots these survival rates. Naturally, the curve for the younger cohort is the one further to the right, reflecting the higher survival rates.

We are also interested in calibrating the difference in lifetime earnings between the two. Using PSID data we find that the difference in average real income between these same cohorts is 28.6 percent. Therefore, we multiply the income profile of the group matched with the young cohort by 1.286, and solve the model for both groups: one with high income and higher life expectancy and the other with lower life-time earnings and lower life expectancy.

Naturally higher income individuals save more. For example, the average age-65 accumulated wealth of higher income males is approximately 40% higher than the one for lower income males. With respect to their retirement decision we have two conflicting effects. On one hand,

²²This is also one of the motivations in the demographic literature for the introduction of cohort effects in the models.

higher income individuals have higher life expectancy which induces them to retire later. However, since they also accumulate more wealth they have an incentive to start enjoying more leisure earlier. We find that, under our calibration, the two effects offset each other almost exactly: the fraction of retirees at age 65 is 18% among high-income individuals and 19% among low-income individuals. More importantly, the high income group benefits more from investing in longevity bonds (Table 7). Since high income individuals live longer on average and also save more, they benefit more (in US dollar terms) from being able to hedge shocks to life-expectancy. As a proportion of age-65 wealth, the benefits of longevity bonds are slightly smaller for the higher income group, due to their higher accumulated savings and later retirement.

6 Security design implications

Currently, there does not exist an active and liquid market for longevity bonds. For that reason it is difficult to parameterize the moments of their return distribution (mean and standard deviation). In this section we consider alternative parameterizations, which also provide important insights into the optimal design of these securities. We also discuss the potential counterparties for these assets.

In order to better understand the benefits of longevity bonds, we consider a higher value for their return volatility, namely $\sigma^L = 10\%$. The resulting asset allocation is shown in Figure 9. Since the bond returns are perfectly correlated with the mortality shocks, a higher volatility

effectively increases the investor's hedging position for a given portfolio allocation. This explains why when we increase σ^L , the portfolio share invested in longevity bonds is, at every age, lower. In a frictionless world this would be a simple re-scaling effect; nothing else would change. However, investors in our model face short-selling constraints, which are binding. The higher return volatility allows them to achieve levels of hedging that were previously unfeasible. This can be seen in Figure 9 by comparing a scaled down version of the baseline portfolio share to that for the higher volatility case.²³ From middle age onwards the agents invest more in longevity bonds than what they would do if they were simply trying to replicate the hedge position for the $\sigma^L = 3\%$ case. The difference is particularly significant as investors approach retirement, and during retirement itself.

The utility gains of investing in longevity bonds are also generally higher when the volatility of returns is equal to 10%. For instance, for females, and for the lower mean improvement and volatility/fixed replacement ratio equal to 0.4 parameterization, the gains of investing in longevity bonds are 2,214 dollars (3.34% of retirement wealth) when $\sigma^L = 10\%$, compared to 1,653 dollars (2.49% of retirement wealth) when $\sigma^L = 3\%$ (Table 4, Panel B). The reason is that the higher volatility allows investors to achieve the same degree of hedging, with a smaller proportion of their wealth allocated to longevity bonds. These are important results to keep in mind for the optimal design of longevity bonds. They show that for a given correlation between

²³This scaled down version is simply the baseline-case portfolio rule, adjusted for the ratio of baseline to new volatility, thus delivering the same exact hedging position.

the returns to longevity bonds and aggregate mortality shocks, it is *typically* better to develop bonds with significant return volatility.²⁴

We use the word *typically* because Figure 9 also reveals an important trade-off: excessive levels of volatility might deter young households from buying the asset. The figure shows that early in life, the portfolio allocation for the $\sigma^L = 10\%$ case is actually slightly lower than the simple scaled-down version of $\sigma^L = 3\%$ case. At this stage of the life-cycle, borrowing constraints and labor income risk are particularly important, and longevity bonds with large return volatility constitute a poor hedge against income risk. This deters households from investing in them, and naturally this crowding out effect increases with the longevity bond return volatility.

It is important to consider not only the effects of return volatility, but also the risk premium on the bonds. For this purpose, an interesting reference is Cairns, Blake and Dowd (2006) who use pricing information contained in the announcement of the European Investment Bank longevity bond that, together with stochastic mortality projections, allows them to estimate the risk premium on new longevity bond issues, as a function of both bond maturity and the cohort index chosen for the floating payments. They point out that their calculations should be viewed as an upper bound on the actual risk premium as they might also be capturing a liquidity risk premium on the EIB bond. For an initial age of cohort equal to 60 (65), and bond

²⁴We could have increased σ^L further, but this would not have had any meaningful impact on the utility gains, since the short-selling constraints are no longer binding for almost any agent in the simulations.

maturity equal to 20 years, the estimated range for the longevity bond risk premium is between 4.8 (12.4) and 8.9 (14.7) basis points.²⁵

Also relevant are the papers that try to determine the value of life insurance-linked assets and liabilities (e.g. Dahl and Moller (2006)). Of particular interest is Cowley and Cummins (2005) who study the bond issued by Swiss Re aiming to insure the company against a large increase in mortality rates (due to, for an example, a pandemic), and the consequent increase in payouts of life insurance policies. The bond payoffs were structured as a call option spread on the mortality index. There was a large demand for this bond. Based on these papers there is no consensus whether there is likely to be more demand pressure from agents wishing to hedge decreases in life expectancies (such as Swiss Re), or from those wishing to hedge increases in life expectancies (such as pension funds). The recent interest of pension plans in longevity swaps suggests that there might be more hedging pressure from the latter, but it hard to give a definite answer.

Therefore we have decided to solve the model for the two possibilities, of a positive and a negative bond risk premium (equal to plus and minus ten basis points), corresponding to scenarios in which there is more hedging pressure from life insurance companies and pension funds, respectively. The welfare gains of investing in longevity bonds are shown in Table 8. In this table we report values both for the Lee-Carter and the high income cohort of the RH model,

²⁵The risk premium decreases significantly with bond maturity, and the lowest maturity that they consider is 20 years. The asset in our model is a one-year bond.

when the replacement ratio is equal to 0.4. Naturally, as we consider a positive (negative) risk premium, the benefits of investing in longevity bonds increase (decrease). Interestingly, for most calibrations the increase in benefits for the positive risk premium case is much larger than the (absolute value of the) decrease that occurs for a negative risk premium. But even for the latter, the benefits of longevity bonds are substantial, and as high as 6 percent of retirement wealth for the higher income cohort.

7 Conclusion

The large somewhat unexpected increases in life expectancy that have occurred over the last few decades have contributed to the current underfunding of defined benefit pension plans. As a result, many of them have either had to reduce promised benefits to employees or to increase contributions. Looking forward, pension plans must decide how to address the issue of longevity risk. Ultimately, they should try to do so if that would be in the interest of their beneficiaries. In this paper we have quantified, in the context of a life-cycle model, the impact of longevity risk on individual saving and retirement decisions, and determined the extent to which individuals would benefit from the hedging of such risk.

We have started by studying the historical evidence and the current projections for life expectancy. The latter show that there is considerable uncertainty with respect to future improvements in mortality rates. We have used this evidence to parameterize a life-cycle model

with longevity risk, in the context of which we have assessed how much longevity risk affects the consumption/saving, retirement and portfolio decisions of an individual saving for both buffer-stock and retirement motives. We have shown that there are several ways in which the agents in our model react to shocks to life-expectancy. First, since longevity risk is realized slowly over the life-cycle, agents optimally save more throughout the life-cycle in response to an improvement in longevity. Second, when faced with large improvements in life-expectancy individuals decide to retire later, even though this entails a utility cost.

We have shown that even when agents optimally decide how much to consume/save and when to retire in response to shocks to life expectancy, they still benefit from being able to invest in financial assets that allow them to insure against longevity risk. More precisely, we have found that when we parameterize longevity risk to match the current projections of the US Social Security and particularly when we try to match those of the UK GAD, the benefits of investing in assets whose returns are correlated with shocks to life expectancy can be economically very significant. They are particularly large in the current context of declining benefits in defined benefit pension plans, and if such declining benefits are a direct response to improvements in life expectancy. These findings lend support to the idea that pension plans should, through their investment decisions, try to hedge the longevity risk that arises from their liabilities, since this would be in the interest of pension plan members.

There seems to be evidence of a link between mortality rates and income, with higher mortality improvements occurring for individuals in the top half of the earnings distribution

(e.g. Waldron, 2007). We have used our model to investigate the effects of such links. We have found that individuals who expect to live longer and who have higher lifetime earnings, save more. The benefits of hedging longevity risk depend on the source of the correlation. If individuals are identical ex-ante, but there is positive correlation between permanent income shocks and mortality improvements, then labor income acts as a hedge for longevity risk. However, we have also shown that if individuals are ex-ante heterogeneous, then individuals who expect to live longer and to have higher lifetime earnings would benefit more from the investment in longevity bonds.

Finally, even though our model is one of the demand for insurance against longevity risk, we have used it to shed some light on the optimal design of longevity bonds, and we have also discussed who might be a counterparty for such bonds. This is an area that deserves further research.

References

- Balduzzi, P., Lynch, A., 1999. Transaction costs and predictability: some utility cost calculations. *Journal of Financial Economics*, 52 (1), 47-78.
- Blake, D., Burrows, W., 2001. Survivor Bonds: Helping to Hedge Mortality Risk. *The Journal of Risk and Insurance*, 68 (2), 339-348.
- Brown, J., Davidoff, T., Diamond, P., 2005. Annuities and Individual Welfare. *The American Economic Review*, (95), 5 1573-1590.
- Brown, J., Poterba, J., 2006. Household Ownership of Variable Annuities. NBER Working Paper, No. 11964.
- Cairns, A., Blake, D., Dowd, K., Coughlan, G., Epstein, D., Khalaf-Allah, M., 2009. A Quantitative Comparison of Stochastic Mortality Models Using Data from England and the United States. Unpublished working paper, The Pensions Institute, Cass Business School.
- Cristia, J., 2007. The Empirical Relationship Between Lifetime Earnings and Mortality. Unpublished working paper, Congressional Budget Office.
- Chai, J., Horneff, W., Maurer, R., Mitchell, O., 2009. Extending Life Cycle Models of Optimal Portfolio Choice: Integrating Flexible Work, Endogenous Retirement, and Investment Decisions with Lifetime Payouts. NBER Working Paper, No. 15079.

- Cagetti, M., DeNardi, M., 2006. Entrepreneurship, Frictions, and Wealth. *Journal of Political Economy*, 114 (5), 835-70.
- Campbell, J., Chan, Y., Viceira, L., 2003. A Multivariate Model of Strategic Asset Allocation. *Journal of Financial Economics*, 67 (1), 41-80.
- Cairns, A., Blake, D., Dowd, K., 2006. A Two-Factor Model for Stochastic Mortality with Parameter Uncertainty: Theory and Calibration. *The Journal of Risk and Insurance*, 73 (4), 687-718.
- Carroll, C., 1997. Buffer-Stock Saving and the Life-Cycle/Permanent Income Hypothesis. *Quarterly Journal of Economics* 114, 433–495.
- Cocco, J., Gomes, F., Maenhout, P., 2005. Portfolio Choice Over The Life-Cycle. *Review of Financial Studies* 18: 491-533.
- Dahl, M., 2004. Stochastic Mortality in Life Insurances: Market Reserves and Mortality-Linked Insurance Contracts. *Insurance: Mathematics and Economics*, 35, 113-136.
- Deaton, A., 1991. Savings and Liquidity Constraints. *Econometrica* 59, 1221–1248.
- Deaton, A., Paxson, C., 2001. Mortality, Education, Income and Inequality Among American Cohorts. In: Wise, D. (Ed.), *Themes in the Economics of Aging*, Chicago University Press for NBER, 129-70.

- Deaton, A., Paxson, C., 2001. Mortality, Income, and Income Inequality in Britain and The United States. NBER working paper 8534.
- Farhi, E., Panageas, S., 2007. Saving and Investing for Early Retirement: A Theoretical Analysis. *Journal of Financial Economics*, 83, 87-212.
- Friedberg, L., Webb, A., 2005. Life is Cheap: Using Mortality Bonds to Hedge Aggregate Mortality Risk. Unpublished working paper, Center for Retirement Research, Boston College.
- Gomes, F., Kotlikoff, L., Viceira, L., 2008. Optimal Life-Cycle Investing with Flexible Labor Supply: A Welfare Analysis of Life-Cycle Funds. NBER Working Paper 13966.
- Gourinchas, P., Parker, J., 2002. Consumption over the life cycle. *Econometrica*, 70, 47-89.
- Government Actuary's Department. National Population Projections: Review of Methodology for Projecting Mortality. National Statistics Quality Review Series Report No. 8.
- Hari, N., Melenberg, B., Nijman, T., Waegenare, A., 2008. Longevity Risk in Portfolios of Pension Annuities. Unpublished working paper, University of Tilburg.
- Human Mortality Database. University of California, Berkeley (USA), and Max Planck Institute for Demographic Research (Germany).
- Available at www.mortality.org.

- Inkmann, J., Lopes, P., Michaelides, A., 2008. How Deep is the Annuity Market Participation Puzzle?. *Review of Financial Studies*, forthcoming.
- Lee, R., Carter, L., 1992, Modelling and Forecasting U.S. Mortality. *Journal of the American Statistical Association*, 87, 419, p. 659-671.
- Lusardi, A., Mitchell, O., 2007. Baby Boomers Retirement Security: The Roles of Planning, Financial Literacy, and Housing Wealth. *Journal of Monetary Economics*, 54, 205-224.
- Menoncin, F., 2008. The Role of Longevity Bonds in Optimal Portfolios. *Insurance: Mathematics and Economics*, Vol. 36, p. 469-484.
- Mitchell, O., Poterba, J., Warshawsky, M., Brown, J., 1999. New Evidence on Money's Worth of Individual Annuities. *The American Economic Review*, 89, 1299-1318.
- Munk, C., Sorensen, C., 2010. Dynamic Asset Allocation with Stochastic Income and Interest Rates. *Journal of Financial Economics*, forthcoming.
- Post, T., 2011. Individual Welfare Gains from Deferred Life-Annuities under Stochastic Mortality. Unpublished working paper, Maastricht University.
- Renshaw, A., Haberman, S., 2006. A Cohort-Based Extension to the Lee-Carter Model for Mortality Reduction Factors. *Insurance: Mathematics and Economics* 38: 556-70.
- Social Security Trustees, 2008. The 2008 OASDI Trustees Report. Section V.A, Demographic assumptions and methods.

- Tauchen, G., Hussey, R., 1991. Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models. *Econometrica*, 59, 371-396.
- Viceira, L., 2001. Optimal Portfolio Choice for Long-Horizon Investors with Nontradable Labor Income. *Journal of Finance*, 55, 1163-98.
- Waldron, H., 2005. Literature Review of Long-Term Mortality Projections. *Social Security Bulletin*, 66, No. 1.
- Waldron, H., 2007. Trends in Mortality Differentials and Life Expectancy for Male Social Security-Covered Workers, by Average Relative Earnings. ORES Working Paper Series, 108, Social Security Administration.
- Yogo, M., 2009. Portfolio Choice in Retirement: Health Risk and the Demand for Annuities, Housing, and Risky Assets. Unpublished working paper, The Wharton School.
- Zeldes, S., 1989, Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence. *Quarterly Journal of Economics* 104, 275-298.

Appendix: A Three Period Model

In order to gain intuition for the asset allocation decision we solve a three period version of our model. The agent maximizes lifetime utility given by:

$$E \left[\frac{C_1^{1-\theta}}{1-\theta} + p_1 \delta \frac{C_2^{1-\theta}}{1-\theta} + p_1 \tilde{p}_2 \delta^2 \frac{C_3^{1-\theta}}{1-\theta} \right], \quad (11)$$

where p_1 and p_2 denote the conditional survival probabilities. In period 1 the individual consumes (C_1) and allocates her savings between riskless bonds and longevity bonds:

$$W_2 = (W_1 - C_1)[\alpha(1 + \tilde{R}_2^L) + (1 - \alpha)(1 + R)] + \tilde{Y}_2, \quad (12)$$

where α denotes the fraction of wealth invested in longevity bonds. The investor's decisions are a function of the current level of wealth (W_1), and depend both on her expected future labor income (Y_2 and Y_3), and on the survival probabilities (p_1 and p_2). Although the individual knows p_1 , she does not know the value of p_2 . Longevity bonds allow her to insure against this uncertainty. In period 2 the individual chooses her optimal consumption (C_2), but now based on the exact realization of the survival probability p_2 . In the last period she simply consumes all available wealth (so that $C_3 = W_3$). We are interested in understanding the behavior of α , for different assumptions about labor income, in order to try to gain intuition for what happens at the different stages of the life-cycle.

In Panel A of Figure A.1 we plot the first-period share invested in longevity bonds, as a function of cash-on-hand, when labor income is constant and riskless: $Y_2 = Y_3 = 10$. This

mimics the retirement phase in our life-cycle model, which is easier to describe. We find that the policy function is decreasing in financial wealth. With constant relative risk aversion households always want to insure a constant fraction of their total wealth. As financial wealth increases, relative to future labor income, this is achieved by decreasing the fraction invested in longevity bonds.

It is important to note that the optimal fraction invested in the bonds is *not* 100%, even though the only source of risk in this version of the model is longevity risk, and longevity bonds provide insurance against such risk. The reason for this perhaps surprising result, is to realize that longevity bonds are an asset that is very different from riskless annuities. When agents invest in longevity bonds they know that this asset yields a low return if p_2 happens to be low. From the perspective of period 1, this implies that the agent will anticipate having less total resources in period 2 in such scenario. Since she cannot off-set this by borrowing against her future labor income then she must trade-off the increase in consumption risk in period 2, against the decrease in conditional consumption risk in period 3.

This trade-off is clearly illustrated in Panel B of Figure A.1, where for comparison we plot the policy function for the case in which $Y_3 = 2 \times Y_2$. In this case period 3 resources are much higher than period 2 resources. Therefore the investor is less interested in transferring additional resources to period 3, and the demand for longevity bonds decreases significantly, except for very high values of wealth.²⁶ Riskless bonds allow the agent to transfer the desired

²⁶The share invested in longevity bonds is higher for very high values of wealth, because there is a lower level

level of wealth from period 1 to period 2 without risk. On the other hand, investing in longevity bonds introduces risk. The compensation for this risk, which is the potential for higher savings for period 3 if p_2 happens to be higher than expected, is of limited value since the agent doesn't want to save for period 3 (Y_3 is much higher than Y_2).

In the third experiment we try to capture investor behavior during working life: second period labor income is now modeled as risky (Y_2 is either equal to 1 or 19, with equal probability). The share invested in longevity bonds is shown in Panel C of Figure A.1. For very low levels of cash-on-hand there is no demand for longevity bonds: labor income risk is much more important. From the perspective of hedging this risk, riskless bonds clearly dominate. It is only when wealth increases that the individual starts to invest in longevity bonds to hedge longevity risk.

of savings (and thus a lower ratio of financial wealth to future labor income).

Table 1: Average annual increases in life expectancy in number of years for a 65 year old individual

Country	United States	Canada	England	Sweden	Germany	France	Italy	Japan
Sample period	1933-2007	1921-2007	1841-2009	1751-2008	1956-2008	1816-2007	1872-2007	1947-2009
Period	Panel A: Males							
Whole sample	0.07	0.06	0.04	0.03	0.10	0.04	0.06	0.14
Pre - 1970	0.02	0.01	0.01	0.01	-0.04	0.01	0.02	0.12
1970 -	0.12	0.12	0.15	0.10	0.14	0.15	0.13	0.16
1960-1969	-0.01	0.03	-0.01	-0.02	-0.08	-0.03	-0.08	0.07
1970-1979	0.13	0.09	0.07	0.04	0.12	0.14	0.08	0.22
1980-1989	0.07	0.08	0.12	0.12	0.13	0.15	0.14	0.16
1990-1999	0.10	0.11	0.15	0.11	0.13	0.11	0.12	0.08
2000 -	0.20	0.23	0.26	0.16	0.19	0.21	0.21	0.18
Period	Panel B: Females							
Whole sample	0.09	0.09	0.05	0.04	0.13	0.06	0.09	0.19
Pre - 1970	0.09	0.07	0.03	0.02	0.07	0.04	0.06	0.15
1970 -	0.09	0.10	0.12	0.11	0.14	0.16	0.16	0.21
1960-1969	0.07	0.14	0.07	0.09	0.03	0.06	0.03	0.12
1970-1979	0.19	0.15	0.08	0.15	0.17	0.19	0.16	0.23
1980-1989	0.03	0.08	0.10	0.12	0.15	0.16	0.17	0.22
1990-1999	0.02	0.06	0.10	0.07	0.12	0.12	0.13	0.19
2000 -	0.15	0.14	0.21	0.10	0.13	0.19	0.17	0.21

Note to Table 1: This table shows average annual increases in life expectancy for a 65 year old individual over time and for different countries. Panel A (Panel B) shows the results for males (females). The data is from the Human Mortality Database.

Table 2.1: Survival Probability Parameters

Description	Parameter	US Historical	US SS Proj.	UK GAD Proj.
Panel A: Male				
Drift	μ_k	-0.648	-0.648	-0.928
Volatility	σ_k	1.861	4.84	10.238
Panel B: Female				
Drift	μ_k	-0.896	-0.496	-0.836
Volatility	σ_k	1.816	5.811	9.624

Note to Table 2.1: This table reports the values for the parameters for the stochastic mortality process for alternative parameterizations. Panel A (Panel B) reports the values for males (females).

Table 2.2: Increase in number of years in age 67 life expectancy, between 2007 and 2045, for different percentiles of the distribution

Percentile of the distribution			
Panel A1: Model Calibration Male			
	Perc 25	Perc 50	Perc 75
US Historical	1.6	2.3	2.9
US Social Security	0.9	2.5	4.0
UK GAD	0.4	3.8	6.6
Panel A2: Data Projections Male			
	Low	Interm.	High
US Social Security	1.0	2.5	4.5
UK GAD	0.6	3.8	8.0
Percentile of the distribution			
Panel B1: Model Calibration Female			
	Perc 25	Perc 50	Perc 75
US Historical	2.7	3.3	3.8
US Social Security	0.5	2.4	4.2
UK GAD	0.6	3.7	6.3
Panel B2: Data Projections Female			
	Low	Interm.	High
US Social Security	0.9	2.4	4.3
UK GAD	0.7	3.7	7.5

Note to table 2.2: The first panel of the table shows the increases in life expectancy at age 67 predicted by the model over a 37 year period for different model parameters and for percentiles 25, 50 and 75 of the distribution of shocks to life expectancy. The second part of each panel shows the projected increases in cohort life expectancy, in number of years, at age 67, between 2007 and 2045, by the US Social Security and by the UK Government Actuary's Department for the low, intermediate and high cost projections. Panel A (Panel B) shows the values for males (females).

Table 3: Other Model Parameters

Description	Parameter	Value
Time Parameters		
Initial age	t_0	30
Min retirement age	t_R^{min}	65
Max retirement age	t_R^{max}	70
Terminal age	T	110
Preference Parameters		
Discount rate	δ	0.98
Risk aversion	θ	3
Preference for leisure	γ	0.9
Bequest motive	b	0
Labor Income		
St dev of temporary income shocks	σ_η	0.0738
St dev of permanent income shocks	σ_u	0.01065
Replacement ratio	λ	0.68212
Asset Returns		
Interest rate	R_f	0.015
Mean Long. Bonds Return	μ^L	0.015
Stdev Long. Bonds Return	σ^L	0.03

Note to Table 3: This table reports the values for the remaining model parameters.

Table 4: Utility Gains of Investing in Longevity Bonds: Alternative Retirement Benefits

Longevity scenarios (vol. and mean improv.)/ Replacement ratio	Repl. Ratio = 0.68		Repl. Ratio = 0.40	
	Age 30 (dollar)	Age 65 (% of wealth)	Age 30 (dollar)	Age 65 (% of wealth)
Panel A: Male				
Historical/Fixed	52	0.29	148	0.24
Lower/Fixed	343	2.00	1025	1.71
Higher/Fixed	947	5.56	2770	4.88
Lower/Stochastic	1193	5.15	1575	2.46
Higher/Stochastic	2468	9.62	3549	6.19
Panel B: Female				
Historical/Fixed	31	0.14	65	0.10
Lower/Fixed	363	1.81	1076	1.61
Higher/Fixed	729	3.53	2155	3.25
Lower/Stochastic	1222	5.35	1653	2.49
Higher/Stochastic	2036	7.68	2874	4.46

Note to table 4: This table reports the utility gains measured at ages 30 (in 2007 US dollars) and at age 65 (as a percentage of average accumulated retirement wealth) for alternative retirement benefits and longevity projections. Panel A (Panel B) shows the results for males (females). The lower (higher) volatility and mean improvement scenario corresponds to the parameter models that better fit the forward looking US Social Security (UK Government Actuary Department) longevity projections.

Table 5: Utility Gains of Investing in Longevity Bonds: Alternative Bequest Parameterizations, Replacement Ratio = 0.4

Longevity scenarios (vol. and mean improv.)/ Replacement ratio	No bequest motive		Joy of giving		Altruistic	
	Age 30 (dollar)	Age 65 (% wealth)	Age 30 (dollar)	Age 65 (% wealth)	Age 30 (dollar)	Age 65 (% wealth)
Panel A: Male						
Lower/Fixed	1025	1.71	737	1.11	1202	1.46
Higher/Fixed	2770	4.88	2324	3.62	3918	4.46
Lower/Stochastic	1575	2.46	1201	1.72	1805	2.13
Higher/Stochastic	3549	6.19	3007	4.82	5434	6.79
Panel B: Female						
Lower/Fixed	1076	1.61	802	1.12	1602	1.89
Higher/Fixed	2155	3.25	1779	2.49	2822	3.36
Lower/Stochastic	1653	2.49	1276	1.79	1920	2.25
Higher/Stochastic	2874	4.46	2396	3.43	3252	3.88

Note to table 5: This table reports the utility gains measured at ages 30 (in 2007 US dollars) and at age 65 (as a percentage of average accumulated retirement wealth) for alternative bequest parameterizations, retirement benefits and longevity projections based on the Lee-Carter model. In the joy of giving parameterization the individual derives utility from the wealth bequeathed. In the altruistic parameterization the individual derives utility from the impact of the wealth bequeathed in the utility of his/her descendent. Panel A (Panel B) shows the results for males (females). The lower (higher) volatility and mean improvement scenario corresponds to the parameter models that better fit the forward looking US Social Security (UK Government Actuary Department) longevity projections.

Table 6: Utility Gains of Investing in Longevity Bonds: Correlated Income and Longevity Shocks.

Longevity scenarios (vol. and mean improv.)/ Replacement ratio	Repl. Ratio = 0.68		Repl. Ratio = 0.40	
	Age 30 (dollar)	Age 65 (% of wealth)	Age 30 (dollar)	Age 65 (% of wealth)
Panel A: Male				
Lower/Fixed	125	0.73	594	1.00
Higher/Fixed	589	3.48	2158	3.83
Lower/Stochastic	757	3.30	1042	1.64
Higher/Stochastic	1907	7.56	2877	5.08
Panel B: Female				
Lower/Fixed	155	0.77	704	1.06
Higher/Fixed	418	2.03	1654	2.53
Lower/Stochastic	794	3.50	1141	1.73
Higher/Stochastic	1516	5.78	2279	3.57

Note to table 6: This table reports the utility gains measured at ages 30 (in 2007 US dollars) and at age 65 (as a percentage of average accumulated retirement wealth) for alternative retirement benefits and longevity projections, and when permanent income shocks are positively correlated with the shocks to improvements in life expectancy, with a correlation coefficient equal to 0.35. Panel A (Panel B) shows the results for males (females). The lower (higher) volatility and mean improvement scenario corresponds to the parameter models that better fit the forward looking US Social Security (UK Government Actuary Department) longevity projections.

Table 7: Utility Gains of Investing in Longevity Bonds: Renshaw Haberman Model, High Income and Low Income Groups.

Longevity scenarios (vol. and mean improv.)/ Replacement ratio	Repl. Ratio = 0.68		Repl. Ratio = 0.40	
	Age 30 (dollar)	Age 65 (% of wealth)	Age 30 (dollar)	Age 65 (% of wealth)
Panel A1: Male, Low Income				
Lower/Fixed	359	2.22	952	1.60
Higher/Fixed	1034	6.08	2663	4.59
Lower/Stochastic	1077	5.22	1459	2.34
Higher/Stochastic	2476	9.28	3311	5.62
Panel A2: Male, High Income				
Lower/Fixed	400	1.76	1119	1.42
Higher/Fixed	1237	5.24	3356	4.36
Lower/Stochastic	1359	4.80	1790	2.17
Higher/Stochastic	3150	8.89	4294	5.53
Panel B1: Female, Low Income				
Lower/Fixed	342	1.89	1573	2.52
Higher/Fixed	719	3.73	2763	4.17
Lower/Stochastic	1504	5.59	2234	3.25
Higher/Stochastic	2532	7.65	3652	5.33
Panel B2: Female, High Income				
Lower/Fixed	434	1.68	1683	1.88
Higher/Fixed	895	3.26	3079	3.39
Lower/Stochastic	1955	5.04	2409	2.57
Higher/Stochastic	3165	6.86	4082	4.41


Note to table 7: This table reports the utility gains measured at ages 30 (in 2007 US dollars) and at age 65 (as a percentage of average accumulated retirement wealth) for alternative retirement benefits and longevity projections, and for different income groups, estimated using the Renshaw Haberman Model. Panel A (Panel B) shows the results for males (females). The lower (higher) volatility and mean improvement scenario corresponds to the parameter models that better fit the forward looking US Social Security (UK Government Actuary Department) longevity projections.

Table 8: Utility Gains of Investing in Longevity Bonds: Alternative Return Parameterizations, Replacement Ratio = 0.4

Longevity scenarios (vol. and mean improv.)/ Replacement ratio	Baseline		Neg. Prem./High Vol		Pos. Prem./High Vol.	
	Age 30 (dollar)	Age 65 (% wealth)	Age 30 (dollar)	Age 65 (% wealth)	Age 30 (dollar)	Age 65 (% wealth)
Panel A1: Male, Lee-Carter Model						
Lower/Fixed	1025	1.71	242	0.40	2915	4.85
Higher/Fixed	2770	4.88	2540	4.47	6813	12.00
Lower/Stochastic	1575	2.46	917	1.43	4192	6.55
Higher/Stochastic	3549	6.19	3404	5.94	8690	15.16
Panel A2: Female, Lee-Carter Model						
Lower/Fixed	1076	1.61	283	0.42	2860	4.29
Higher/Fixed	2155	3.25	1380	2.08	5067	7.64
Lower/Stochastic	1653	2.49	932	1.40	4261	6.42
Higher/Stochastic	2874	4.46	2680	4.15	6892	10.69
Panel B1: Male, RH Model, High Income						
Lower/Fixed	1119	1.42	201	0.26	3347	4.24
Higher/Fixed	3356	4.36	3020	3.92	8269	10.74
Lower/Stochastic	1790	2.17	931	1.13	4911	5.96
Higher/Stochastic	4294	5.53	4936	6.36	10543	13.58
Panel B2: Female, RH Model, High Income						
Lower/Fixed	1683	1.88	553	0.62	4061	4.53
Higher/Fixed	3079	3.39	1996	2.20	6909	7.60
Lower/Stochastic	2409	2.57	1189	1.27	5592	5.97
Higher/Stochastic	4082	4.41	3361	3.63	9051	9.78


Note to table 8: This table reports the utility gains measured at ages 30 (in 2007 US dollars) and at age 65 (as a percentage of average accumulated retirement wealth) for alternative return parameterizations, retirement benefits and longevity projections. Panel A shows the results for the Lee-Carter model and Panel B shows the results for the high income group using the Renshaw Haberman model. The neg. prem./high vol. shows the results for as specification with lower mean return and higher volatility. The pos. prem./high vol. shows the results for as specification with higher mean return but also higher volatility. The lower (higher) volatility and mean improvement scenario corresponds to the parameter models that better fit the forward looking US Social Security (UK Government Actuary Department) longevity projections.

Figure 1: Life expectancy in the United States at selected ages



Note to Figure 1: This figure shows period life expectancy over time and at selected ages (birth, age 30, and age 65) for the United States for males and females. The data is from the Human Mortality Database from 1933 to 2007.


Figure 2: Conditional probability of death for a male US individual




Note to Figure 2: This figure shows the conditional probability of death over the life-cycle for selected years (1960, 1970, 1980, 1990, 2000 and 2007) and for a male United States individual. The data is from the Human Mortality Database.

Figure 3, Panel A: Projected increases in life expectancy for a 65 year old US and UK individual

Panel A: Male




Panel B: Female



Note to Figure 3: This figure plots the projected increases in cohort life expectancy, in number of years, over time for a 65 year old male (Panel A) and female (Panel B) US and UK individual. This figure plots a principal, a high and a low variant. The projections were done by the US Social Security (2009 OASDI Trustees Report) and by the UK Government Actuaries Department (www.gad.gov.uk).


Figure 4: Actual and estimated conditional survival probabilities




Note to Figure 4: This figure shows the data and the estimated conditional probabilities of death for a United States male individual at selected years. The data used in the estimation is from the Human Mortality Database from 1933 to 2007. The estimation is done using the Lee-Carter model.

Figure 5: Model implied age 67 life-expectancy for different parameters of the $k(x)$ stochastic process and the Lee-Carter model

Panel A: Male




Panel B: Female



Note to Figure 5: This figure plots the model implied life expectancy at age 67 for different parameters of the stochastic process for $k(x)$ and the probability that such expectancy will occur. Panel A (Panel B) plots the results for males (females).

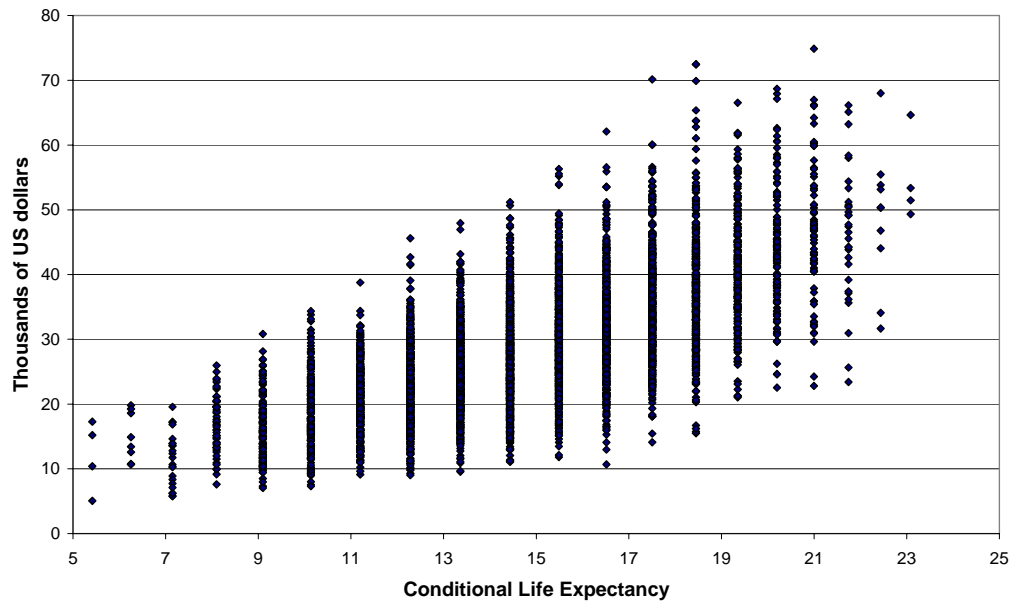
Figure 6: Model Results

Panel A: Simulated Consumption, Income, Wealth, and Retirement Decisions




Note to Figure 6, Panel A: This figure plots the simulated profiles over the life-cycle in the baseline model. The figure plots an average across 5,000 simulated profiles, and the values are converted into 2007 US dollars.

Panel B: Age 65 wealth as a function of age 65 life expectancy




Note to Figure 6, Panel B: This figure plots the simulated wealth accumulation (in 2007 US dollars) at age 65 for 5000 simulated agents, in the baseline model, as a function of their life-expectancy (at age 65).

Figure 7: Asset Allocation to Longevity Bonds




Note to Figure 7: This figure plots the average allocation to longevity bonds across 5000 simulated agents, in the baseline version of the model (late in life wealth accumulation is equal to zero hence the portfolio decision is not identified)

Figure 8: Survival Rates for Different Cohorts in the RH Model



Note to Figure 8: This figure plots the survival rates for two different cohorts estimated using the RH model. The data is from the HMD for the US from 1933 to 2007.


Figure 9: Simulated Portfolio Allocation to Longevity Bonds for different values of longevity bonds return volatility.




Note to Figure 9: This figure plots the simulated share of wealth invested in longevity bonds for two different parameterizations of the model, with different levels of longevity bonds return volatility: 3% and 10%. In addition the figure also plots the allocation for the 3% volatility case, re-scaled by 3/10 thus corresponding to the same hedge position as the allocation for the 10% volatility case. The figure plots an average across 5,000 simulated profiles.

Figure A.1: Portfolio Allocation to Longevity Bonds in the 3-period model, as function of wealth.


Panel A: Baseline Case



Panel B: Different Levels of Period 3-Income



Panel C: The Impact of Income Risk



Note to Figure A.1: This figure plots the portfolio share invested in longevity bonds as a function of cash-on-hand in the three-period model for different assumptions regarding income and its risk.