

Longevity Risk and Retirement Savings

João F. Cocco¹

Francisco J. Gomes²

This version: March 2009

¹Universidade Católica Portuguesa, Palma de Cima, 1649-023 Lisboa, Portugal, CEPR, and Nestpar (on leave from London Business School). Tel (21) 721 4000. Email jcocco@fcee.ucp.pt.

²London Business School, Regent's Park, London NW1 4SA, UK and CEPR. Tel (020) 7000 8215. Email fgomes@london.edu. We would like to thank David Laibson, Bertrand Melenberg, Alex Michaelides, Olivia Mitchell, Pedro Santa-Clara, and seminar participants at 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.

Abstract

Over the last couple of decades there have been unprecedented, and to some extent unexpected, increases in life expectancy which have raised important questions for retirement savings. We study optimal consumption and saving choices in a life-cycle model, in which we allow for changes in the distribution of survival probabilities, according to the Lee-Carter (1992) model. We allow individuals to hedge longevity risk through an endogenous retirement decision and by investing in financial assets designed to hedge this risk. We find that the welfare gains of investing in these assets are substantial when longevity risk is calibrated to match the forward-looking projections of the US Social Security Administration and the UK Government Actuaries Department. These gains are particularly large in a context of declining benefits in defined benefit pension plans. 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 13.04 years.³ Three and a half decades later, in 2005, a 65-year-old male had a life expectancy of 17.2 years. This represents an increase of 1.16 years per decade. To understand what such an increase implies in terms of the savings needed to finance 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 \$10.52 in 1970, to \$13.53 by 2005. This is an increase of roughly 29 percent. Or in other words, to finance a given stream of real consumption during retirement, a 65-year-old male would have needed 29 percent more wealth in 2005 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 actuaries and insurers. This is hardly surprising given the historical evidence on life expectancy. From 1970 to 2005 the average increase in the life expectancy of a 65 year old male was 1.12 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 of a 65 year old male was 1.35 years/decade from 1970 to 2005, but only 0.17 years/decade from 1870 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 company sponsored pension plans. For 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.

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, and have instead chosen to contribute towards personal pensions that tend to be defined contribution in nature. This means a transfer of longevity risk from pension providers

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

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

to the individuals themselves.

It is true that the risk that the individual might live longer than average may be reduced by the purchase of annuities at retirement age. However, a young individual saving for retirement faces substantial uncertainty as to the level of aggregate life expectancy, and consequently the annuity prices, that she will face when she retires.⁵ Furthermore, markets may be incomplete in the sense that individuals may lack the financial assets that would allow them to insure against this risk. 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.

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 Lee-Carter model (1992), which is the leading statistical model of mortality in the demographic literature.

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 they are stochastic. Naturally the individual forms an expectation of such probabilities when making her decisions. In addition to allowing the individual to adjust her savings in response to changes in life expectancy, we allow for endogenous retirement, i.e. we allow the individual to choose the age at which to retire.

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, with mixed results. For example, in December 2003, Swiss Re. issued a \$400m three-year life catastrophe bond.⁶ This was a direct attempt by Swiss Re. to insure itself against a catastrophic mortality deterioration (e.g. a

⁵In this respect our paper differs from the large and important literature on annuities. See for example the recent contributions of Mitchell, Poterba, Warshawsky, and Brown (1999), Brown, Davidoff, and Diamond (2005), Brown and Poterba (2006), and Inkmann, Lopes and Michaelides (2008).

⁶Blake 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.

pandemic). This bond offered an opposite hedge to pension funds and other annuity providers. In spite of this and other attempts, a liquid market for the trading of longevity risk does not currently exist.

It is with this process of financial innovation in mind, and to investigate the role of such bonds in individual portfolios, that we allow the agent in our model to invest in longevity bonds, or financial assets whose returns are correlated with the shocks to the survival probabilities. We study portfolio allocation between these bonds and riskfree assets, how it 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.⁷

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 an optimal response to improvements in life expectancy. This is the case even though such decision carries a utility cost. 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.

However, and importantly, our model shows 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 welfare implications. More precisely, we show that individuals in our model 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 welfare gains are substantially higher when the payouts of defined benefit pension plans are lower, and when they are negatively correlated with aggregate survival rates.

⁷There 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 this paper we take bond prices as given and investigate their role in household portfolios. Menoncin (2007) also introduces longevity bonds in an optimal savings and portfolio choice problem but in his model there is no labor income or retirement.

This scenario is motivated by recent events, which suggest that as survival rates increase, retirement benefits are progressively decreased. 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.

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.

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 sections 3 and 4 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 5. In section 6 we consider a version of the model with misperceptions. These misperceptions are motivated by the fact that official life tables are sometimes mistakenly interpreted by users as allowing for subsequent mortality changes. The final section concludes.

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. It is important to note that period life tables are sometimes mistakenly interpreted by users as allowing for subsequent mortality changes.

We focus our analysis on life expectancy at ages 30 and 65. 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. Our 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 individual for the United States and England over time, at birth, age 30, and at age 65.

Table 1 reports average annual increases in life expectancy for a 65 year old male 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 increases in life expectancy since 1970, and that these increases have not been confined to the US or England. Furthermore, except for Japan, there does not seem to be evidence that the increases in life 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 0.11 years. In England, the corresponding values are 0.07 and 0.15. These increases 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 from this figure the probability of death has decreased substantially from 1970 to 2000, 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 2005, so that to facilitate comparison we use the projections provided in the 2004 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.

Figure 3 plots the projected future increase, in number of years, in the cohort life expectancy of a 65 year old male, relative to 2005.⁸ 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 a lower expected increase in life expectancy, and a lower dispersion in its estimates than the UK GAD. For instance, the UK GAD projects an intermediate (low/high) increase in life expectancy between 2005 and 2040 of 3.3 years (0.4/7) whereas the corresponding numbers for the US are 2.2 years (0.8/3.9). 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.

⁸Cohort 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.

3 A Model of Longevity Risk

3.1 Survival Probabilities

We solve a life-cycle model of consumption and savings, in the spirit Carroll (1997) and Gourinchas and Parker (2001), but in which survival probabilities are stochastic. We let t 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 Lee-Carter (1992) model to describe survival probabilities. This is the leading statistical mortality model in the demographic literature, and it has been shown to fit the data relatively well. In addition, it has the advantage of being a relatively simple model. Mortality rates are given by:

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

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.

Lee and Carter (1992) show that a random walk with drift describes 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.

3.2 Preferences

Let p_t denote the probability that the individual is alive at age $t + 1$, conditional on being alive at age t , so that $p_t = 1 - m_t$. For a given individual age and time are perfectly co-linear, so that in order to simplify the exposition 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:

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

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. 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.

3.3 Labor Income

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

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

where $f(t, Z_t)$ is a deterministic function of age and of a vector of other individual characteristics, Z_t , ε_t is an idiosyncratic temporary shock distributed as $N(0, \sigma_\varepsilon^2)$, 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 ε_t . 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. After this age income is modeled as a constant fraction

λ of permanent labor income in the last working-year:

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

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} .

3.4 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 are longevity bonds, a financial asset 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:

$$R_t^L = \mu^L + \frac{\sigma^L}{\sigma^k} \varepsilon_t^k$$

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) shock to the mortality rates.

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.

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

$$W_{t+1} = (1 + R_{t+1}^p)(W_t + Y_t - C_t). \quad (6)$$

where

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

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.

4 Calibration

4.1 Time and 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.

4.2 Survival probabilities

Undoubtedly, the calibration of the parameters for the mortality process are 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 order to parameterize the stochastic process for survival probabilities we do two things. 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 reflect historical data, other information, and expectations of future improvements in mortality.

For the estimation of the Lee-Carter model using historical data, we use US data from 1933 to 2005, which is the data period available in the Human Mortality Database, and estimate:

$$\ln(m_{t,x}) = a_t + b_t \times k_x + \varepsilon_{t,x} \quad (7)$$

where $\varepsilon_{t,x}$ is an error term with mean zero and variance σ_ε^2 , which reflects particular age-specific historical influences not captured by the model. 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 2000. 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. In Figure 5 we plot the evolution over time of the k_x parameter. This figure confirms, from a different perspective, the data shown in Table 1. After 1970 there has been a significant decrease in k_x reflecting the decreases in mortality rates that have taken place since then.

We use the time series data of k_x to estimate the parameters of the random walk. The estimated

drift parameter μ^k is -0.736 and the standard deviation of the shocks σ^k is 1.716 (both of these parameters are reported in Table 2). 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 2005 (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 65 in 2005 compare to life expectancy at age 65 when the individual reaches such age (i.e. in year 2040)? In other words, we ask which is the increase in the life expectancy of a 65 year old individual that the model forecasts to take place over the next 35 years?

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 35 years. Therefore, in Table 3 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 3 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 2.25 years. The 25th and 75th percentiles are 1.62 and 2.86 years, respectively.

The calculations in the first row of Table 3 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 2005, 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 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 the second panel of table 3, and based on the same data that we have used in figure 3, we report the US Social Security and the UK GAD projected increases in the life expectancy of a 65 year old over a 35 year period (from 2005 to 2040). 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 the first row of Table 3), to the projections of the US Social Security and the UK GAD (shown in the second panel of table 3) we see that the former lead to lower dispersion in the forecasts, particularly so when compared to the projections of the UK GAD.⁹ To motivate this 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 have experimented with different drift and volatility parameters for the Lee-Carter model, and projected simulated mortality improvements under such alternative parameterizations. In the second row of the first panel of table 3 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 slightly smaller drift parameter and a higher volatility than those estimated from historical data. The reduction in drift and the increase in volatility are much larger if we wish to roughly match the UK GAD projections instead (third row of the first panel of table 3). We solve our model both for the parameters estimated using historical data and for these alternative parameterizations. As a baseline case we use the parameters that better match the US Social Security projections.

4.3 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, which are also reported in Table 2.

⁹One 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.

Deaton and Paxson (2001) have investigated the correlation between aggregate labor income and mortality improvements using UK and US data and concluded that the two aggregate series do not seem to be correlated. Therefore we assume zero correlation between labor income shocks and shocks to longevity. 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 these bonds are a fairly recent financial development. The returns on longevity bonds are linked to an aggregate mortality index, but there is not a perfect (or even linear) relationship between the two. 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.

On the other hand, the longevity bond proposed by the European Investment Bank in November 2004 would have had floating coupons that would be directly linked to a mortality index. Using mortality rate forecasts and the Lee-Carter model Friedberg and Webb (2005) compute the hypothetical returns for this longevity bond if it had been available over a long period. They obtain a return volatility of 3%, which we use as the baseline value in our calibration ($\sigma^L = 3\%$). They also find that the returns are negatively correlated with aggregate consumption and, in a CCAPM framework with relative risk aversion of 10, the implied risk premium would be -0.02% . Therefore, we set our baseline risk premium to zero (that is $\mu^L = 1.5\%$).

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.

4.4 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.

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.

5 Results

5.1 Life-Cycle Profiles

We first solve the model assuming that the individual does not have access to longevity bonds. These bonds are a recent phenomenon and they are not widely available. We use the optimal policy functions to simulate the consumption and savings profiles of five thousand agents over the life-cycle. In Figure 7 we plot the average simulated income, wealth and consumption profiles. 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, Figure 7 plots the proportion of individuals working. This proportion declines from age 65 onwards, and only a small proportion of individuals are still working 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 large improvements in life-expectancy to those who do not face such improvements. 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 Table 4 we report simulated age 65 variables, namely average consumption, cash-on-hand, and retirement decisions, for different realizations of the shocks to life expectancy. From this table 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 accumulated higher levels of savings. Thus, in our model, the 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).

Individuals in our model can also decide at which age to retire. In the last column of Table 4 we report the proportion of individuals aged 65 who decide to postpone retirement, i.e. who decide to keep on working. The vast majority of individuals with low life expectancy decide to retire at age 65, but those individuals who have longer life expectancy tend to keep on working. 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.

Finally, Table 4 shows that individuals with longer life expectancy not only have higher accumulated savings and retire later, but they also tend to consume less. This is of course due to the fact that they have a higher expected number of years of consumption to finance out of current savings and income.

5.2 Financial Innovation: Longevity Bonds

We now consider the role of financial innovation, by allowing the individual in our model to invest in longevity bonds, in addition to the risk-free asset. The welfare gains associated with the investment in these bonds also provide us with a measure of the extent to which the individuals in our model are affected by longevity risk. We first develop a three period model, to gain intuition on the effects at work, before turning our attention to the more realistic full life-cycle model.

5.2.1 Three period model

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

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

where p_1 and p_2 denote the conditional survival probabilities. In period 1 the individual con-

sumes (C_1) and allocates her savings between riskless bonds and longevity bonds:

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

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 8 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 8, 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.¹⁰ Riskless bonds allow the agent to transfer the desired 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). These policy functions will allow us to explain the behavior of the portfolio allocation during the retirement phase of our model.

In the third experiment that we consider 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 8. 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.

5.2.2 Life-cycle results

With this intuition in mind we turn our attention to the more realistic and empirically parameterized life-cycle model. Figure 9 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. By definition, the riskless asset is more suited to hedge this risk and, following the intuition described in panel C of Figure 8 (in the 3-period model), these households prefer to invest a significant fraction of their wealth in riskless bonds. However, as the agent accumulates more wealth, and starts saving for retirement, the allocation to longevity bonds quickly increases, and reaches 100%. During retirement the portfolio allocation to longevity bonds remains high, but it soon becomes undetermined since household wealth converges to zero.

¹⁰The share invested in longevity bonds is higher for very high values of wealth, because there is a lower level of savings (and thus a lower ratio of financial wealth to future labor income).

5.3 Welfare Gains

In this section we compute the welfare 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 welfare gains under the form of consumption equivalent variations. 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 welfare 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 use the riskfree rate to calculate the present value of this difference.¹¹

It is important to point out that the welfare measure that we calculate is an ex-ante welfare measure, in the sense that it is a welfare measure calculated as of age 30, for a given longevity risk. More precisely we compare the welfare 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 welfare gains as of age 30, we also calculate the welfare gains, as of age 65, as a fraction of accumulated retirement wealth. This number is obtained simply by multiplying the age 30 welfare gains by $(1 + R_f)^{35}$, i.e. capitalizing the welfare gains to age 65, and dividing 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 welfare gains is simply a re-scaling of the former.

The welfare results are shown in Table 5. For the parameters that match the forward-looking US SS projections, the welfare gains of investing in longevity bonds are 144 US dollars measured

¹¹Alternatively, we could have calculated the welfare losses in terms of the percentage gain/loss in the consumption stream.

at age 30, or 1.22 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. For such parameters the welfare gains are 412 dollars as of age 30, or as high as 3.23 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 when we allow for endogenous retirement and optimal saving and consumption decisions as a response to longevity shocks.

5.3.1 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. This means that in the future the level of benefits that individuals will derive from defined benefit schemes are likely to be smaller than the ones that we have estimated using historical data. This is important since defined benefit pension plans, because of their nature, provide insurance against longevity risk. Thus one might argue that for a high level of defined benefit pensions, and from the point of view of individuals covered by such pensions, longevity risk does not matter much, but given the current and expected reductions in the level of such benefits, longevity risk will become more important.

We use our model to investigate the extent to which that is likely to be the case. 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 on welfare. The results are shown in Panel B of table 5. The age 30 welfare gains associated with the longevity bonds are considerably higher, and equal to 436 and 1129 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 shown in Panel A. 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 from having access to longevity bonds are an order of magnitude higher than before, and equal to 3.59 percent and 5.85 percent of age 65 retirement wealth, for the US SS and the UK GAD projections scenarios, respectively. These are very significant gains.

5.4 Security Design Implications

We now consider an alternative calibration of the volatility of the shocks to the longevity bond return. More precisely, we set $\sigma^L = 10\%$. This analysis provides important insights into the optimal design of these securities.

The results for portfolio allocation are shown in Figure 10. Since the returns on the longevity bonds 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 the longevity bonds is lower, at every age. In a frictionless world this would be a simple re-scaling effect, and nothing else should change. However, in our model the investor faces short-selling constraints, which are binding. The higher volatility allows her to achieve levels of hedging that were previously unfeasible. This can be seen by comparing a scaled down version of the baseline portfolio share, to the portfolio share for the higher volatility case.¹² From middle age onwards the agent invests more in longevity

¹²The scaled down version is simply the old portfolio rule, adjusted for the ratio of old to new volatility: thus delivering the same exact hedging position.

bonds than what she would do if she was simply trying to replicate the hedge position for the $\sigma^L = 3\%$ case. The difference is particularly significant as she approaches retirement, and during retirement itself.

This is also reflected in the welfare gains of investing in longevity bonds, which are mostly higher in this case (Panel C of Table 5). This is because for the case of a higher volatility the agent can with a smaller proportion of his wealth allocated to longevity bonds achieve the same degree of hedging. These are important results to keep in mind for the optimal design of longevity bonds. More precisely, our results show that for a given correlation between the returns to longevity bonds and aggregate mortality shocks, it is *typically* better to develop bonds with significant return volatility.¹³ We use the word *typically* because Figure 10 also reveals an important trade-off: excessive levels of volatility might deter young households from buying this asset. 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. This, is due to the borrowing constraints and labor income risk. At this stage of the life cycle, households are less willing to buy longevity bonds because they constitute a poor hedge against labor income risk. As we increase their return volatility this crowding out effect becomes stronger. These are important considerations to keep in mind for the optimal design of longevity bonds.

One of the difficulties associated with developing a market for longevity risk is that there is not a natural counterparty to supply the bonds. Our model sheds some light on this issue. More precisely, in Panel D of Table 5 we report the welfare gains associated with the investment in longevity bonds where there is a bequest motive, with the parameter b set equal to one. Comparing the welfare gains to those reported in panel A, we see that the welfare gains from investing in longevity bonds are now smaller. These results show that individuals with a bequest motive may be in a good position to provide insurance against longevity risk to those who do not have a bequest motive. This is intuitive: individuals with a strong bequest motive have a longer horizon (possibly infinite), and for this reason they are not as affected by longevity risk.

¹³We could have increased σ^L further, but this would not have had any meaningful impact on the welfare gains, since the short-selling constraints are no longer binding for almost any agent in the simulations.

6 Model with Misperceptions

In the model in the previous section, agents were fully rational in the sense that they were at each point in time aware of the increases in longevity, and rationally took into account any expected future increases in longevity. We now allow for misperceptions on the part of the agent. We motivate these misperceptions in the following way. Consider an individual who uses official period life tables to make his consumption and saving decisions, without realizing that these are “only” *period* life tables. That is, we investigate the consumption and saving choices of an agent who in each period looks up mortality rates in official period life tables, and in this way is informed about survival rates, but who fails to recognize that such survival rates are likely to improve in the future. This is a common mistake made by users of life tables, who think these official life tables allow for future mortality improvements. In terms of our model this means that the agent thinks that the drift parameter for the stochastic process for longevity (μ^k) is zero when making his consumption/saving decisions.

We compute two types of welfare gains. First, we compute the welfare gains or losses associated with the mistakes. That is: the welfare loss relative to the case in which the individual does not make such mistakes, without having access to longevity bonds. Naturally, these welfare gains are computed using the objective probability measure, and thus reflect the *real* welfare gains that the typical investor will realize. The results are shown in Panel A of Table 6. It shows that, as of age 30, the individual loses 701 dollars for failing to recognize future mortality improvements in his decisions. These welfare losses are considerably larger, equal to 1728 dollars, if we consider instead the parameters calibrated using the GAD projections. The welfare losses are also economically very significant when calculated as a fraction of retirement wealth: 5.94 and 13.53 percent for the US and UK parameterizations, respectively. The reason for these very large welfare losses is the fact that the agent, by failing to anticipate future mortality improvements saves significantly less.

The case in which individuals fail to anticipate future mortality improvements, even though they are aware of current life expectancy, may be considered too extreme. Therefore, in panel B of table 6 we consider an alternative scenario, in which agents think that $\mu^k = -0.736$, the value that we have estimated from US historical data, when in reality the values are those implied by the US SS and the UK GAD projections, reported in Table 2. Naturally, the losses associated with the mistakes are smaller, but they still are economically significant, particularly

for the UK parameterization.

Panel C of table 6 reports the welfare gains from having access to longevity bonds in the case of the same misperceptions as in Panel B. Interestingly, the gains of having access to longevity bonds are smaller than the ones obtained when the agent is fully rational (previously reported in the first panel of Table 5). The reason is that in the case of misperceptions the agent thinks that longevity bonds are not as useful, and decides to invest less in these bonds. Therefore, although the gains from *actually investing* in longevity bonds would still be large, the gains from *being able to invest* in them are actually small.

The results in Table 6 show how important it is for agents to be informed about current life expectancy, and also future mortality improvements. In particular, for the market for longevity bonds to function properly, it is important that investors do have a good understanding of life tables and in particular of the limitations of official (period) life tables.

7 Conclusion

In this paper we have documented the existing empirical evidence, and the current projections on life expectancy. 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 agent in our model reacts to shocks to life-expectancy. First, since longevity risk is realized slowly over the life-cycle, the agent, throughout the life-cycle, optimally saves more in response to an improvement in longevity. Second, when faced with large improvements in life-expectancy the individual decides to retire later, even though this entails a utility cost.

We have shown that even when the agent optimally decides how much to consume/save and when to retire as response to shocks to life expectancy, she benefits from being able to invest in financial assets that allow her 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 longevity bonds can be economically very significant. They are particularly large in a context of declining benefits in defined benefit pension plans, and if such declining benefits are a direct

response to improvements in life expectancy.

We have also shown that when agents are informed about life expectancy, but make an incorrect assessment of the probability of future improvements in life expectancy, the effects of such mistakes on individual welfare can be substantial. This is a common mistake since official life tables tend to be period life tables, which do not allow for future mortality improvements. Finally, our paper has shed some light on the optimal design of longevity bonds, and who might be a counterparty for such bonds.

References

- Blake, David and William Burrows, 2001, “Survivor Bonds: Helping to Hedge Mortality Risk”, *The Journal of Risk and Insurance*, 68 (2), 339-348.
- Brown, Jeffrey, Thomas Davidoff, and Peter Diamond, 2005, “Annuities and Individual Welfare,” *The American Economic Review*, (95), 5 1573-1590.
- Brown, Jeffrey, and James Poterba, 2006, “Household Ownership of Variable Annuities,” *NBER Working Paper*, No. 11964.
- Cagetti, Marco and Mariacristina DeNardi, 2006, “Entrepreneurship, Frictions, and Wealth,” *Journal of Political Economy*, 114 (5), 835-70.
- Cairns, Andrew, David Blake and Kevin Dowd, 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, Christopher D., 1997, “Buffer-Stock Saving and the Life-Cycle/Permanent Income Hypothesis,” *Quarterly Journal of Economics* 114, 433–495.
- Cocco, Joao, Francisco Gomes, and Pascal Maenhout, 2005, “ Portfolio Choice Over The Life-Cycle,” *Review of Financial Studies* 18: 491-533.
- Dahl, Mikkel, 2004, “Stochastic Mortality in Life Insurances: Market Reserves and Mortality-Linked Insurance Contracts,” *Insurance: Mathematics and Economics*, 35, 113-136.
- Deaton, Angus S., 1991, “Savings and Liquidity Constraints,” *Econometrica* 59, 1221–1248.
- Deaton, Angus S. and Christina Paxson, 2001, “Mortality, Income, and Income Inequality in Britain and The United States” *NBER working paper* 8534.
- Friedberg, Leora and Anthony Webb, 2005, “Life is Cheap: Using Mortality Bonds to Hedge Aggregate Mortality Risk,” Working Paper, Center for Retirement Research, Boston College.
- Gomes, Francisco, Laurence Kotlikoff, and Luis Viceira, 2008, “Optimal Life-Cycle Investing with Flexible Labor Supply: A Welfare Analysis of Life-Cycle Funds,” *NBER Working Paper* 13966.

- Gourinchas, Pierre-Olivier, and Jonathan Parker, 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
- Human Mortality Database. University of California, Berkeley (USA), and Max Planck Institute for Demographic Research (Germany).
Available at www.mortality.org.
- Inkmann, Joachim, Paula Lopes, and Alexander Michaelides, 2008, "How Deep is the Annuity Market Participation Puzzle?," *Working Paper*, London School of Economics.
- Lee, R.D. and L.R. Carter, 1992, "Modelling and Forecasting U.S. Mortality," *Journal of the American Statistical Association*, 87, 419, p. 659-671.
- Lusardi, Annamaria and Olivia S. Mitchell, 2007, "Baby Boomers Retirement Security: The Roles of Planning, Financial Literacy, and Housing Wealth," *Journal of Monetary Economics*, 54, 205-224.
- Menoncin, Francesco, 2008, "The Role of Longevity Bonds in Optimal Portfolios," *Insurance: Mathematics and Economics*, Vol. 36, p. 469-484.
- Mitchell, Olivia, James Poterba, Mark Warshawsky, and Jeffrey Brown, 1999, "New Evidence on Money's Worth of Individual Annuities," *The American Economic Review*, 89, 1299-1318.
- Hari, Norbert, Bertrand Melenberg, Theo Nijman, and Anja De Waegenaere, 2008, "Longevity Risk in Portfolios of Pension Annuities," *Working Paper*, University of Tilburg.
- Social Security Trustees, 2008, "The 2008 OASDI Trustees Report," Section V.A, Demographic assumptions and methods.
- Tauchen, G. and R. Hussey, 1991, "Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models," *Econometrica*, 59, 371-396.

Waldron, Hilary, 2005, “Literature Review of Long-Term Mortality Projections,” *Social Security Bulletin*, 66, No. 1.

Zeldes, Stephen, 1989, “Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence”, *Quarterly Journal of Economics* 104, 275–298.

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

	United States	Canada	England	Sweeden	Germany	France	Italy	Japan
Sample Period	1959 - 2005	1921 - 2003	1841 - 2005	1751 - 2004	1956 - 2002	1899 - 2003	1872 - 2003	1947-2004
Whole Sample	0.07	0.05	0.04	0.03	0.09	0.06	0.05	0.15
Pre 1970	0.02	0.01	0.01	0.01	-0.04	0.03	0.03	0.12
1970 -2000	0.11	0.09	0.12	0.09	0.13	0.13	0.11	0.15
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	0.08	0.08	0.12	0.12	0.13	0.15	0.14	0.16
1990 - 1999	0.11	0.11	0.15	0.11	0.14	0.11	0.12	0.08

Note to Table 1: This table shows average annual increases in life expectancy for a 65 year old male over time and for different countries. The data is from the Human Mortality Database.

Table 2: Parameters of the model

Description	Parameter	Value
Survival probabilities: US Historical		
Drift	μ_k	-0.736
Volatility	σ_k	1.716
Survival probabilities: US Social Security Projections		
Drift	μ_k	-0.756
Volatility	σ_k	4.033
Survival probabilities: GAD Projections		
Drift	μ_k	-1.256
Volatility	σ_k	8.067
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_ϵ	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 2: This table reports the parameters of the model.

Table 3: Increase in number of years in age 65 life expectancy, between 2005 and 2040, for different percentiles of the distribution

Model	Percentile of the distribution		
	25	50	75
Historical: $\sigma_k = 1.716$, $\mu_k = -0.736$	1.62	2.25	2.86
US: $\sigma_k = 4.033$, $\mu_k = -0.756$	0.60	2.10	3.54
UK: $\sigma_k = 8.066$, $\mu_k = -1.256$	0.34	3.31	5.99
Data	Low	Intermediate	High
US Social Security Projections	0.8	2.1	3.7
UK GAD Projections	0.4	3.3	7

Note to table 3: The first panel of the table shows the increases in life expectancy at age 65 predicted by the model over a 35 year period for different model parameters and for percentiles 25, 50 and 75 of the distribution of shocks to life expectancy. The second panel of the table shows the projected increases in cohort life expectancy, in number of years, at age 65, between 2005 and 2040, by the US Social Security and by the UK Government Actuary's Department for the low, intermediate and high cost projections.

Table 4: Model simulated age 65 cash-on-hand, consumption, and retirement decisions, for different realizations of the shocks to life expectancy

Life Expectancy	Cash-on-hand	Consumption	Prop. working
75.7	36.4	26.8	0.00
77.2	40.1	26.8	0.00
78.7	40.4	26.4	0.21
80.3	40.7	26.1	0.41
81.9	42.1	26.0	0.74
83.4	44.3	26.0	0.91
84.8	45.7	25.9	0.93
86.2	47.8	25.9	0.95
87.5	49.8	25.9	0.99
88.6	50.1	25.7	1.00
89.7	51.9	25.7	1.00

Note to table 4: This table reports age 65 simulated cash-on-hand, consumption and proportion of individuals still working, for different realizations of the shocks to life expectancy. The parameters of the model are set at their benchmark levels. The results are averages across 5,000 individuals.

Table 5: Welfare Gains of Investing in Longevity Bonds

	Welfare gain	
	At age 30 (Dollar value)	At age 65 (% of Ret. Wealth)
Panel A: Baseline		
US Social Security Projections	143.78	1.22
UK GAD Projections	412.06	3.23
US Historical	25.63	0.21
Panel B: Alternative retirement benefits		
US SS Proj. Lower Replacement Ratio	435.74	1.04
US SS Proj. Stochastic Replacement Ratio	623.91	3.59
UK GAD Proj. Lower Replacement Ratio	1128.89	2.69
UK GAD Proj. Stochastic Replacement Ratio	1341.90	5.85
Panel C: Higher Return Volatility		
US SS Proj.	135.59	1.15
US SS Proj. Lower Repl. Ratio	438.87	1.05
US SS Proj. Stochastic Repl. Ratio	908.42	5.23
UK GAD Proj.	494.60	3.87
UK GAD Proj. Lower Repl. Ratio	1444.90	3.35
UK GAD Proj. Stochastic Repl. Ratio	2430.96	10.60
Panel D: Bequest motive		
US SS Proj.	49.75	0.20
UK GAD Proj	215.40	0.85


Note to table 5: This table reports the welfare gains measured at ages 30 and at age 65 (as a percentage of average accumulated retirement wealth) for different parameters of the model. Panel A shows the results for the baseline parameterizations, panel B for alternative retirement benefits, and panel C for the case of a higher standard deviation of the returns on the longevity bond, equal to 10%, both for the baseline and for the alternative retirement benefits scenarios. Panel D reports results for the case in which the bequest motive parameter is set equal to one.

Table 6: Welfare Gains: Model with Misperceptions

	Welfare gain	
	At age 30 (Dollar value)	At age 65 (% of Ret. Wealth)
Panel A: Failing to anticipate future improvements		
US Social Security Projections	700.98	5.94
UK GAD Projections	1727.55	13.53
Panel B: Underestimating the magnitude of future improvements		
US Social Security Projections	130.33	1.11
UK GAD Projections	819.91	6.42
Panel C: Invest. in Long. Bonds: Under. the magnitude of fut. impr.		
US Social Security Projections	100.22	0.85
UK GAD Projections	201.62	1.58

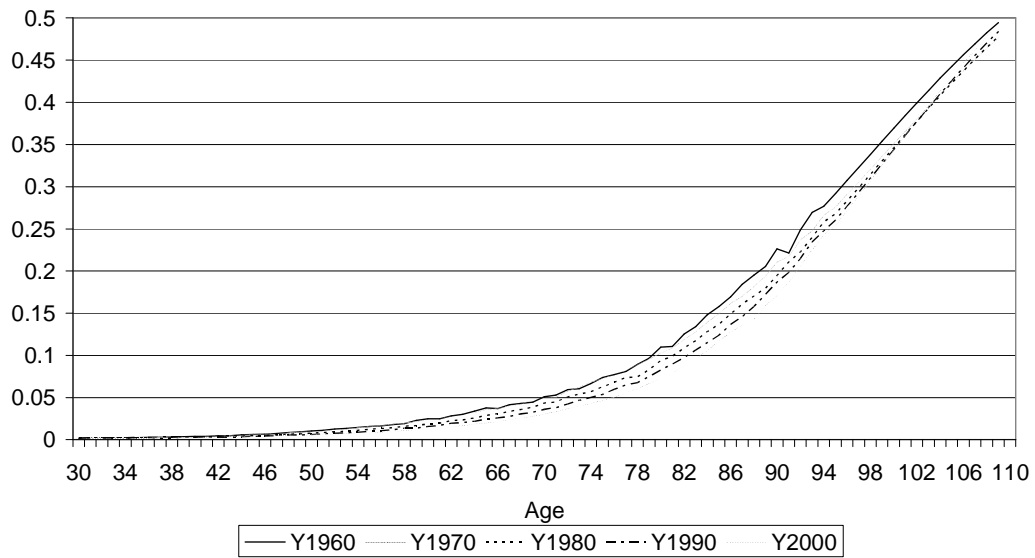
Note to table 6: This table reports the welfare gains measured at ages 30 and at age 65 (as a percentage of average accumulated retirement wealth) for different parameters of the model with misperceptions. Panel A shows the welfare loss for an agent who fails to anticipate future mortality improvements. Panel B shows the welfare loss for an agent who thinks that future mortality improvements are similar to the historical average. Panel C shows the welfare gains from investing in longevity bonds for an agent who thinks that future mortality improvements are similar to the historical average.

Figure 1: Life expectancy in the United States and England for a male individual 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 and for England. The data is from the Human Mortality Database. The data for the United States is from 1933 to 2005, and for England is from 1841 to 2005.

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, and 2000) and for a male United States individual. The data is from the Human Mortality Database.

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




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 US and UK individual. This figure plots a principal, a high and a low variant. The projections were done by the US Social Security (2004 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 2005. The estimation is done using the Lee-Carter model.

Figure 5: Estimated $k(x)$ parameter in the Lee-Carter model



Note to Figure 5: This figure shows the estimated $k(x)$ parameter in the Lee-Carter model. The data used in the estimation is for a male US individual, from the Human Mortality Database from 1933 to 2005.

Figure 6: Model implied age 65 life-expectancy for different parameters of the $k(x)$ stochastic process



Note to Figure 6: This figure plots the model implied life expectancy at age 65 for different parameters of the stochastic process for $k(x)$ and the probability that such expectancy will occur.


Figure 7: Simulated Consumption, Income, Wealth and Retirement Decisions in the Baseline Model




Note to Figure 7: This figure plots the simulated consumption, wealth, and income in the baseline model. The figure plots an average across 5,000 simulated profiles.

Figure 8: 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 8: 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.

Figure 9: Simulated Portfolio Allocation to Longevity Bonds in the Baseline Model



Note to Figure 9: This figure plots the simulated share of wealth invested in longevity bonds in the baseline model. The figure plots an average across 5,000 simulated profiles.

Figure 10: Simulated Portfolio Allocation to Longevity Bonds for different values of longevity bonds return volatility.



Note to Figure 10: 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.