

Mortality Change, the Uncertainty Effect, and Retirement ^{*}

Sebnem Kalemli-Ozcan

David N. Weil

University of Houston

Brown University and NBER

December 2001

Abstract

We examine the role of changing mortality in explaining the rise of retirement over the course of the 20th century. We construct a model in which individuals make labor/leisure choices over their lifetimes subject to uncertainty about their date of death. In an environment in which mortality is high, an individual who saved up for retirement would face a high risk of dying before he could enjoy his planned leisure. In this case, the optimal plan is for people to work until they die. As mortality falls, however, it becomes optimal to plan, and save for, retirement. We simulate our model using actual changes in the US life table over the last century, and show that this “uncertainty effect” of declining mortality would have more than outweighed the “horizon effect” by which rising life expectancy would have led to later retirement. One of our key results is that continuous changes in mortality can lead to discontinuous changes in retirement behavior.

JEL Classification: E21, I12, J11, J26

^{*}We thank Andrew Foster, Herschel Grossman, Robin Lumsdaine, Enrico Spolaore, and the participants at seminars at The Bank of Israel, Bar-Ilan, Ben-Gurion University, Bilkent University, Brown University, Hebrew University, University of Houston, Koc University, University of Maryland, Rice University, University of Rochester, and Tel Aviv University for helpful comments.

1 Introduction

One of the most dramatic economic changes that has taken place in the last 100 years has been the rise in retirement as an important stage of life. At the beginning of the Twentieth century, retirement was a rarity. Many people didn't live into old age, and most of those who did continued to work until shortly before death. By the end of the century, the vast majority of workers could expect to experience a prolonged period of healthy leisure after their working years were over.

Needless to say, the economic repercussions of this change in the life cycle pattern of labor supply have been enormous. Since a significant fraction of consumption during retirement is publicly funded, the growth in retirement has strained government budgets - a phenomenon which will soon be exacerbated by population aging (Weil, 1997). Meanwhile private funding of anticipated retirements has led to the accumulation of vast pools of capital. Exactly what fraction of current capital accumulation can be attributed to life cycle savings is an issue of contentious debate. Modigliani's (1986) estimate is that as much as 80% of wealth can be attributed to life cycle saving, while Kotlikoff and Summers (1988) estimates that it is only 20%. According to Lee (1998), the fraction of wealth attributable to life cycle saving in the US doubled between 1900 and 1990.

Explaining the rise in retirement has been a major endeavor for economists in the last several decades. Three prominent explanations can be discerned. The first is that the public pension programs, such as Social Security in the United States, have been instrumental in pushing workers out of the labor force, particularly through high implicit rates of taxation on wage income earned at older ages (Gruber and Wise, 1998). The second explanation is that rising lifetime income has led workers to optimally choose a larger period of leisure at the end of life (Costa, 1998). The final explanation is that changes in the technology of production have lowered the productivity of older workers, leading employers to seek to get rid of them (Graebner, 1980; Sala-i-Martin, 1996). None of these explanations, either taken separately or as a group, has been a completely satisfactory explanation for the rise in retirement. This paper proposes a new explanation for the rise in retirement, which we call the "uncertainty effect." It is not meant to be a complete substitute for the factors discussed above, but rather as an addition to the set of usual suspects that must be considered in explaining the rise of retirement.

In our model, the driving force behind the change in the life-cycle pattern of labor supply is a change in the pattern of mortality. Individuals make labor/leisure choices over their

lifetimes subject to uncertainty about their date of death. In an environment in which mortality is high, an individual who saved up for retirement would face a high risk of dying before he could enjoy his planned leisure. In this case, the optimal plan is for individuals to work until they die. As mortality falls, however, individuals will find it optimal to plan, and save for, retirement. The effect of falling mortality on labor supply complements several other effects of falling mortality, most notably on human capital investment and fertility, that have been analyzed in recent literature.¹

The rest of our paper is structured as follows. In section two, we briefly examine historical data on mortality and retirement, and also discuss at greater length some of the other explanations for the rise of retirement. Section three presents and solves our model of endogenous retirement in a stripped-down economic environment, in which the other factors affecting retirement that were discussed above are not present. Section four looks at the interaction between our new channel and one of the existing explanations of the rise in retirement. Section five concludes.

2 Historical Data

2.1 Mortality

Falling mortality has been one of the most significant aspects of the process of economic growth over the last several centuries. Male Life expectancy at birth in the United States rose from 40.17 in 1850 to 47.81 in 1900 and 78.42 in 1990 (Haines (1994), Keyfitz and Flieger (1990)). The most significant component of mortality decline has been the reduction in infant and child deaths. Obviously, this aspect of the mortality decline is not related to the issues of retirement saving that we discuss in this paper. Although the decline in adult mortality has not been quite as dramatic as that for children, it has nevertheless be a significant part of the story of economic growth over the last century or more. Figure 1 shows the number of survivors from a cohort of 20-year-olds who would be alive at different ages, using life tables from the United States for 1850, 1900, 1950 and 1990. In 1850, the probability that a 20-year-old would reach age 65 was only roughly 40%. By 1990, it was roughly 80%.²Figure 2 shows the probability of death on a log scale.

¹See Meltzer (1992), Ehrlich and Lui (1991), Eckstein et al. (1998), Kalemli-Ozcan (2001).

²The total number of expected remaining years of life for a 20-year-old male rose from 38 to 52 over this period.

2.2 Retirement

Paralleling the reduction in adult mortality has been a massive increase in retirement. In 1930, the labor force participation rate for men aged 65 and over was 58%. By 2001, it had fallen to 17.5%.³ The trend in retirement prior to the Great Depression has been a subject of controversy among historians. Ransom and Sutch (1986, 1988) claim that labor force participation rate for men over 60 was roughly constant between 1870 and 1930, while Costa (1998) and Moen (1994) argue that labor force participation for older men had been declining since the late Nineteenth Century.⁴

Examining the labor force participation of the elderly does not tell the full story, of course, since, as shown above, a large fraction of the population never lived to this age. Carter and Sutch (1996) estimate that at the beginning of the twentieth century, a 55-year-old man had only a 21.5% probability of retiring before he died, excluding “death-bed retirement” associated with illness in the last weeks of life.

Another way to demonstrate the dramatic rise in retirement is to look at how the expected number of years that an individual would spend retired has changed over time. This measure incorporates both changes in mortality and changes in labor force behavior. The expected fraction of adult life spent retired rose from 6.7% for the cohort of men who began working in 1850 to 12.3% for the cohort that began its working life in 1900 and to an estimated 31.0% for the cohort that began working in 1990 (Lee, 1996).

While we will argue that the decline in mortality is one explanation for the fall in labor force participation of the elderly, it would be foolish to argue that it is the only one. Since we cannot incorporate all of these effects into a single model, our approach will be to ask how large a change in retirement could plausibly be explained by the mortality effect alone, in a model where the other factors are not present. In the rest of this section, however, we briefly discuss three other channels.

The most obvious alternative explanation for the decline in labor force participation of the elderly is the creation of old-age insurance programs like Social Security. Economists differ over the question of what fraction of the increase in retirement can be explained by changes

³Such a change could theoretically be due to a changing age distribution of the population over 65 in the presence of constant age-specific participation rates. This is not the case, however: age specific participation rates have also fallen dramatically. See Costa (1998) and Lumsdaine and Wise (1994).

⁴Matthews, et. al. (1982) give the labor force participation rates for men aged 65 and over in the U.K. declining over time in manner similar to those in the U.S. 1881: 73.6%; 1901: 61.4%; 1921: 58.9%; 1931: 47.9%; 1951: 31.1%; 1973: 18.6%.

in Social Security policies. Danziger, et. al. (1981), reviewing several studies, conclude that it is best to “attribute as an upper bound one-half of the decline in the older male labor force since 1950 to Social Security’s work disincentives.”⁵ By contrast, Gruber and Wise (1998) argue that Social Security and other government policies explain a large fraction of the rise in early retirement. Even if this latter conclusion is correct, however, such an observation still leaves open the question of why these policies have evolved as they did. It may be that Social Security policies themselves should be viewed as an endogenous response to other factors that increased the desirability of retirement.

A second explanation for the rise in retirement focuses on the income effect of higher wages. In a taste-based model of retirement, the association of higher income per capita with longer retirement comes about because people in richer countries choose to spend a higher fraction of their income on retirement. When the wage increases, there is an income effect, which leads to the purchase of more of all normal goods, including leisure, and a substitution effect, which works to reduce leisure by raising its opportunity cost. For higher income to lead to longer retirement, not only must the income effect dominate, but there must be some reason why this leisure must be taken at the end of life, rather than spread out evenly. The best argument for this phenomenon is that there is some kind of non-convexity in the enjoyment of leisure; a simple example would be that only people with full time leisure can move to Florida.⁶

A third explanation for the increase in retirement looks toward production technology. This explanation starts with the notion that older workers are less able to do certain tasks than younger ones. If the nature of production shifts toward those tasks over time, then the labor force participation of older people should go down. The argument that the nature of production has changed is presented in Graebner (1980) and Moen (1988). Moen argues that the decline of agriculture as a source of employment led to an increasing physical separation of the home and the workplace, making a gradual withdrawal from the labor force no longer possible.⁷ Similarly, as factory labor and hourly wages replaced piece rate work, it became

⁵p. 996.

⁶This is the argument of Costa (1998). Similarly, Fields and Mitchell (1984) argue that in the U.S. setting, retirement is well modeled as “a choice based on balancing the monetary gains from continued work versus leisure forgone,” rather than as stemming from mandatory retirement regulations or ill health. Mandatory retirement covered only a minority of workers during the period at which they look, and they estimate that it was binding in the retirement decisions of only two or three percent of retiring workers. The Age Discrimination in Employment Act of 1978 and subsequent legislation, which have virtually eliminated mandatory retirement, have had little effect on actual retirement patterns.

⁷Long (1958) rejects the argument that urbanization was the cause in the decline of labor force participation

harder for an older employee to work at his own pace. Graebner argues that it is the nature of large scale, technologically advanced production to demand standardized workers.⁸

As an alternative to changes in productive technology, the cause of the increase in retirement could theoretically lie in changes in the technology of health and longevity. If the fraction of older people who are unable to perform relevant tasks increases, labor force participation rates should go down. The evidence, however, points in exactly the opposite direction: as retirement has increased the health of the elderly has improved (Fields and Mitchell, 1984). In 1994, 89 percent of those between 65 and 74 reported “no disability whatsoever.”⁹

3 Model

The explanation for rising retirement considered here looks to the very reduction in mortality discussed above. Specifically, we focus on the change in the uncertainty regarding mortality. This point can be seen most clearly by looking at Figure 3. The figure shows the probability of dying at different ages, conditional on having reached age 20, using the cross-sectional male life tables for 1900 and 1990. The fact that the mean age of death goes up (from 61.2 to 73.5) is hardly surprising. What is more interesting is that the standard deviation of the age of death falls, from 18.1 to 14.9. And of course the coefficient of variation in the age of death (the standard deviation divided by the mean) falls by even more, from 0.30 to 0.20. By 2050, the mean age of death is expected to further rise, to 76.5, while the standard deviation will also rise, to 15.4. The coefficient of variation will remain almost exactly constant.

How should declining mortality affect the retirement age? Two different forces are at work. First, and most intuitively, longer life would be expected to increase the number of years during which an individual plans to work, simply because longer life means that there will be more years of consumption which need to be paid for. We call this channel the

for men over 65 over the period 1890-1950. The rate of labor force participation was higher in rural areas, but the decline was sharp in both areas: 5.4% per decade in rural areas, 4.4% per decade in urban areas. The rate of decline for the country as a whole was 5.4% per decade, indicating that there was some effect of movement from rural to urban areas.

⁸A different mechanism by which technological change can lead to retirement is by making old workers obsolete, either because older workers are slower at learning new productive techniques, or because it is not worthwhile for older workers to learn new techniques, since they will have less time in the labor force in which to employ them. Graebner (1980) documents the widespread notion that retirement resulted because it was easier to train new workers than to retrain old ones. Both of these arguments suggest that retirement should be related not to the level of technology, but to its rate of change.

⁹For more on this point see Costa (1998).

“horizon effect” of increased life expectancy. But a second effect arises from the fact that increases in life expectancy may reduce the uncertainty surrounding whether a person will live into old age. Where mortality is high, individuals are unlikely to live into old age, and so any savings for a planned retirement would likely be wasted. The optimal program in such circumstances would be to plan continue working in old age, should such an eventuality arise. In this case, a reduction in mortality, by making survival into old age more likely, can reduce the planned age of retirement. We call this second channel the “uncertainty effect.”

Whether the uncertainty effect or the horizon effect of increased life expectancy is more important depends on the manner in which life expectancy increases, as well as on the nature of the individual’s optimization problem. After setting up the general problem in section 3.1, we make this point in sections 3.2 and 3.3 by considering two different forms of life-expectancy increase, which produce opposite effects on the retirement age. Finally, in section 3.4, we use actual data to assess which effect should have been dominant given the historical change in the life table.

3.1 Setup of the Problem

Consider an individual who receives utility from both consumption and leisure. We assume that there are only two possible levels of labor supply: a person may either be fully employed or retired. We also assume that once a person retires, he may not re-enter the labor force.¹⁰ Let γ represent the increment to utility from leisure due to retirement (that is, the difference between utility from leisure before retirement and after). For convenience, utility from leisure and consumption are taken to be separable. An important feature of this utility function is that marginal utility of retirement leisure does not decline with the length of retirement – that is, there are no decreasing returns to retirement.¹¹¹²

The instantaneous utility function is,

¹⁰The question of why retirement usually takes the form of a sudden and complete withdrawal from the labor force is a difficult one, and we do not address it here.

¹¹This does not mean that utility from consumption and leisure are not treated symmetrically. There may be decreasing returns to leisure within a given day, but this does not effect the retirement decision since the length of the workday is taken as fixed. Allowing for decreasing returns to the period of retirement would be analogous to writing down a utility function in which there were decreasing returns to total lifetime (as opposed to instantaneous) consumption, so that if the discount rate were zero, doubling the lifespan and keeping annual consumption constant would not double utility from consumption.

¹²We also do not consider the possibility that the utility from leisure, or the dis-utility from work, rises with age. This would be a further explanation for why leisure is concentrated at the end of life.

$$\begin{aligned}
U &= \ln(c) && \text{if working} \\
&= \ln(c) + \gamma && \text{if retired.}
\end{aligned} \tag{1}$$

The assumption that utility from consumption is given by the log function is made so that changes in wages will not, by themselves, affect the optimal age of retirement. In general, changes in wages (holding interest rates and mortality constant) will have both income and substitution effects on the demand for end-of-life leisure. With log utility these just balance each other, and so the derivative of the optimal retirement age with respect to the wage is zero. If utility is more curved than the log function, then the income effect will dominate, and increases in wages will lead to a reduction in the optimal retirement age. We explore this effect further in section 4.

Future utility is discounted at rate θ . Thus we can write lifetime utility as,

$$U = \int_0^T e^{-\theta x} [\ln(c(x))] dx + \int_R^T e^{-\theta x} [\gamma] dx, \tag{2}$$

where R is the age of retirement, T is the age of death, and $R \leq T$.

In the case where the date of death is uncertain, expected lifetime utility is given by

$$E(U) = \int_0^\infty e^{-\theta x} P(x) [\ln(c(x))] dx + \int_R^\infty e^{-\theta x} P(x) [\gamma] dx, \tag{3}$$

where $P(x)$ is the probability of being alive at age x . Note that in this case, R is the planned age of retirement, but an individual who dies before age R will not experience any retirement at all. Since the only uncertainty that we admit to the model is about the date of death, and since this uncertainty is not resolved until it is too late to do anything about it, individuals will form time-consistent plans for consumption and retirement at the beginning of their lives.

Individuals who are working receive a wage of w . The real interest rate is r . Assets of an individual aged z are

$$\int_0^{\text{Min}(z,R)} w e^{r(z-x)} dx - \int_0^z c(x) e^{r(z-x)} dx. \tag{4}$$

We impose the condition that an individual cannot die in debt. In the case where the date of death is known, this reduces to the familiar condition that terminal assets are zero. In the case where the date of death is uncertain, by contrast, this condition means that assets at all points in time must be non-negative. In other words, uncertain mortality imposes a liquidity constraint that prevents the individual from borrowing.

3.2 Optimal Retirement with No Uncertainty

We begin by considering the case where the date of death, T , is known with certainty. The individual's lifetime budget constraint is

$$\int_0^R e^{-rx} [w] dx = \int_0^T e^{-rx} [c(x)] dx. \quad (5)$$

The individual will maximize his utility subject to his budget constraint by choosing a path of lifetime consumption, $c(x)$ and an endogenous retirement age, R . The first order condition with respect to consumption implies,

$$\dot{c}(x) = [r - \theta]c(x) \quad (6)$$

The equation above together with the budget constraint will give us the initial consumption, c_0 ,

$$c_0 = \frac{w\theta (1 - e^{-rR})}{r (1 - e^{-\theta T})}. \quad (7)$$

The first order condition with respect to the endogenous retirement age with equation (7) give us an implicit equation for the optimal retirement age, R ,

$$\frac{r(1 - e^{-\theta T})}{\gamma\theta} = e^{(r-\theta)R}(1 - e^{-rR}). \quad (8)$$

The above equation will obviously only hold true if $R \leq T$. The other alternative is that the individual is at a corner solution, where $R = T$.

The derivative of R with respect to T is

$$\frac{dR}{dT} = \frac{r}{\gamma e^{\theta(T-R)}[\theta + (r - \theta)e^{rR}]}. \quad (9)$$

It can be shown that this derivative is positive if optimum $R < T$. Thus in the case where there is no uncertainty, increases in life expectancy will lead to increases in the retirement age.¹³

The dashed line in figure 4 shows how increases in life expectancy affect retirement for a particular set of parameters. The figure shows that the retirement age rises with life expectancy.¹⁴

¹³The derivative in equation 9 is positive since the necessary condition for the second order condition to maximization to hold is that $\theta > (r - \theta)e^{rR}$.

¹⁴In figure 4 we used the following values for the parameters: $r = 0.06, \theta = 0.03, \gamma = 1$.

3.3 Optimal Retirement with Uncertainty and No Liquidity Constraints

We now turn to the case where the date of death is uncertain. Improvements in mortality take the form of changes in age-specific death probabilities.

We begin by considering a form of uncertainty that can be handled tractably in an analytic model. We then turn to a more general consideration of uncertainty, in which we must use a dynamic programming algorithm to find optimal retirement ages and consumption paths.

To introduce uncertainty, suppose the individual has a constant probability ρ of dying. The probability of being alive at age x is $P(x) = e^{-\rho x}$. Thus expected utility at the beginning of life is,¹⁵

$$V = \int_0^\infty e^{-(\theta+\rho)x} [\ln(c(x))] dx + \int_R^\infty e^{-(\theta+\rho)x} [\gamma] dx. \quad (10)$$

As mentioned above, uncertainty about the date of death imposes a liquidity constraint, so that assets at all times must be non-negative. This in turn means that many individuals will die holding positive wealth (although not necessarily, since an individual may choose to work until he dies and hold no assets). Thus accidental bequests will be generated. We ignore these bequests, assuming in effect that they are thrown away.¹⁶

In general optimization problems with liquidity constraints cannot be solved by using standard analytic techniques. In this section we make the model tractable by assuming that $r > \rho + \theta$. This guarantees that individuals will have rising consumption paths over the course of their lifetimes, and thus will hold positive assets at all times. In other words, we choose parameters such that the liquidity constraint never binds. In the next section, where liquidity constraints are considered explicitly, we can relax this constraint on parameters.

We can write the individual's lifetime budget constraint as,

$$\int_0^R e^{-rx} [w] dx = \int_0^\infty e^{-rx} [c(x)] dx. \quad (11)$$

The individual will maximize his expected utility subject to his budget constraint by choosing a path of lifetime consumption, $c(x)$, and an endogenous retirement age, R . The first order condition with respect to consumption implies,

¹⁵Throughout this paper, we are concerned only with how mortality affects retirement and saving decisions, so we think of "birth" as being the beginning of working life. When we consider actual life tables, we begin our analysis assuming that the individual has already survived to age 20

¹⁶If individuals place positive value on bequests, but still value them less than their own consumption, then our main qualitative results will continue to hold.

$$\dot{c}(x) = [r - \theta - \rho]c(x) \quad (12)$$

The equation above together with the budget constraint will give us the initial consumption, c_0 ,

$$c_0 = \frac{w(\theta + \rho)}{r} \left[1 - e^{-rR} \right]. \quad (13)$$

The first order condition with respect to the endogenous retirement age together with the consumption will give us an implicit equation for the optimal retirement age, R ,

$$\frac{r}{\gamma(\theta + \rho)} = e^{(r-\theta-\rho)R} \left[(1 - e^{-rR}) \right]. \quad (14)$$

As in the certainty case, this equation for retirement age will hold only in the case where retirement age is not at a corner solution. In this case the relevant corner solution is that $R = \infty$, in other words, the individual plans never to retire.

We can also calculate the derivative of R with respect to ρ ,

$$\frac{dR}{d\rho} = - \frac{-Re^{(r-\rho-\theta)R} + Re^{-(\rho+\theta)R} + r/(\gamma\theta + \rho)^2}{(r - \rho - \theta)e^{(r-\rho-\theta)R} + (\rho + \theta)e^{-(\rho+\theta)R}}. \quad (15)$$

It can be shown that this derivative is positive for large values of R . Thus, the initial effect of reductions in mortality (starting from a high level of ρ) will be to reduce the age of retirement. However, once retirement age has fallen, it is possible that further improvements in mortality can raise the age of retirement.

This partial equilibrium setup is sufficient to show the endogenous emergence of life cycle saving. For sufficiently high rates of mortality, individuals will find it optimal to set the planned date of retirement, R , to infinity: in other words, they will simply plan to work until they die. The solid line in figure 4 show the relationship between retirement and life expectancy (which is just $1/\rho$) for a particular set of parameters. Notice that R is the *planned* age of retirement, but than many people will not live long enough to reach it. Thus there is no inconsistency in having the retirement age be greater than life expectancy. Figure 4 shows the key result: as life expectancy rises, the planned age of retirement falls!

Figure 4 also allows for an explicit analysis of the role of uncertainty in affecting retirement. In the case where the date of death is known (the dashed line), increases in life expectancy affect retirement only through the “horizon” effect. By contrast, in the case where increases in life expectancy take place due to a falling probability of death, such changes in

life expectancy will be associated with both a change in the agent's horizon and a change in the uncertainty surrounding life expectancy. Thus changes in retirement in this case (the solid line) will reflect both the “horizon” and “uncertainty” effects. The difference between the solid and dashed line, then, is the pure uncertainty effect. It can be seen in figure 4 that for low life expectancies, the uncertainty effect is large, and so the retirement age is much higher under uncertainty than under certainty. As life expectancy rises, however, the gap between these two diminishes – indeed, it is precisely this effect that allows the retirement age in the case of uncertain death to fall even as life expectancy is increasing.

3.4 Realistic Mortality Rates

The analysis in the last section shows that for a simple specification of mortality, increases in life expectancy can actually cause a *decline* in the age of retirement. In other words, it is theoretically possible that the uncertainty effect will dominate the horizon effect. A natural question to ask is whether such an outcome is consistent with actual changes in mortality that have been experienced historically. It is to this question that we now turn.

We use a discrete time version of our model to calculate optimal retirement ages for actual mortality data. We examine life tables for US Males, going back to 1850 and forecasted through 2050. Our procedure for finding the optimal retirement age is straightforward. We loop through possible retirement ages, and for each one calculate optimal consumption and asset paths subject to the constraint that assets are never negative.¹⁷ We then calculated expected utility associated with each possible retirement age.

¹⁷Let P_t be the probability that an individual will be alive in period t , and let A_t be assets at the beginning of time t . The first-order conditions for consumption are

$$c_{t+1} = c_t \frac{(1+r)}{(1+\theta)} \frac{P_{t+1}}{P_t} \quad \text{if} \quad A_{t+1} > 0,$$

$$c_t = w_t \quad \text{if} \quad A_{t+1} = 0.$$

If assets are being carried into period $t + 1$, then it must be the case that the usual first order condition for marginal utilities (adjusted by interest rate, time discount, and mortality probabilities) must hold. The only case in which the condition will not hold is if the individual would like to shift more consumption into period t , but is unable to do so because of the liquidity constraint. In this case he will consume all of his wages.

In the case of the problem being addressed here, we can take advantage of a special feature of mortality rates that always holds true in the data once one moves beyond childhood: mortality rates are an increasing function of age. This delivers the result that, as long as assets are positive, the growth rate of consumption must be declining over time. The only time when consumption growth will not be declining is when assets are zero, in which case consumption is constant and equal to the wages. See Carroll (1997) for a similar dynamic programming algorithm.

Figure 5 shows an example of our calculation of expected lifetime utility as a function of planned retirement age, using mortality data from 1900 through 1990. The optimal retirement age is the age at which this function reaches its maximum. The two interesting points about the picture are that for some of the mortality data the function has two peaks, and that over time the position of one peak versus the other peak changes. This is our main result which is the emergence of retirement. In 1900s the optimal plan was never to retire since lifetime utility is monotonically rising as a function of retirement age all the way to the end of life. By contrast, in 1990 the optimal retirement age is 57. In fact in 1980 the optimal plan changes from never retiring to retiring at age 57.¹⁸

It is interesting to note that the two possible retirement ages in year 1980, which yield very similar levels of expected utility, involve different strategies for lifetime consumption and asset accumulation. These are shown in Figures 6-9. The optimal consumption and asset paths when the expected retirement age is 57 (Figures 6 and 7) have the standard life-cycle shape. The growth rate of consumption is highest at the beginning of life (age 20, in our problem), because the probability of death, which functions like the discount rate, is lowest.¹⁹ Consumption growth falls over time, reaching zero at age 65, then becoming negative. Assets grow from the beginning of life until retirement at age 57, and fall monotonically thereafter.

In the case where the individual plans never to retire, the path for consumption (figure 8) has the same shape (although a higher level) as when retirement age is 57, until the age of 83. After age 83, consumption becomes equal to the wage since liquidity constraint is binding. The path of assets is different than the previous case in two ways (figure 9): First, in the case where there is no retirement, assets reach zero late in life, while in the case of retirement assets in old age are always positive. The second difference between the asset paths when the retirement age varies is in their size: When retirement is at age 57, total assets peak at roughly 20 times annual income; when the individual plans never to retire, assets peak at 9 times annual income. (In the case where there is no retirement, the accumulation of assets is due to the difference between the interest rate and the time discount rate.)

The important implication of figure 5 is that because of the twin peaks in the expected utility function, small changes in the parameters will lead to large changes in expected retirement age. The twin peaks structure of utility in the figure is not a result of any choice of parameters on our part. Rather it results from the dichotomy between the two strategies for optimal consumption mentioned above: plan for a retirement, with the risk that one might

¹⁸The γ values are calibrated as explained below. The other parameters are $r = 0.06, \theta = 0.03$.

¹⁹See Hurd (1990).

die early and thus have wasted all of the money saved up; or plan to pay for consumption in the event that one lives into old age by working. There is no way to “convexify” between these two strategies.

We have conducted robustness checks using different values of the interest rate and the time discount rate. We consistently got the result that reductions in mortality lead to a discrete jump downward, rather than a gradual reduction, in the optimal retirement age. Further, as a test of whether this two-humped pattern was somehow related to peculiarities in the exact data in the life table, we re-ran our analysis using a more structured version of the mortality data. Specifically, for each year, we estimated a Gompertz specification in which the log of the probability of death was regressed on a constant and on age, using data starting at age 20. We then used the fitted values from these regressions as data in our optimization program. The resulting figure looked very similar to the one that we present here.

Taken literally, this two-peaked pattern would imply that there should be a sudden shift in behavior from never retiring to planning for a large retirement. Obviously this is not what is observed in the data, but we do not think of this as a major failing of the model. In the real world, heterogeneity, institutions, learning, and a host of other factors would tend to cause retirement ages to adjust slowly, rather than jumping all at once, in response to a change in mortality.

The optimal retirement age for different life tables and for different values of parameters are given in table 1. We calibrate the value of γ , the parameter that determines the utility of leisure in retirement. Specifically, for each set of values for r and θ , we find the value of γ which makes the optimal retirement age using the 1980 life table equal to 65. (Because of the twin-peakedness of the optimal retirement function, it is often impossible to find a value of γ which will yield an optimal retirement age of exactly 65. In these cases, we choose the value of γ that will yield the largest retirement age that is lower than 65.) We then calculate optimal retirement ages for life tables from other years. The emergence of retirement in response to falling mortality as shown in Figure 5, is also evident here for different parameter values.

4 Income vs Uncertainty Effects on Retirement

In the model as set up in the previous section we suppressed any income effects on retirement by using the log utility function. One natural extension of the model is to allow for income effects, and to see how income effects interact with uncertainty. Technically, this is simply

a matter of replacing the log utility function with a Constant Relative Risk Aversion utility function. As long as the coefficient of relative risk aversion is greater than one, increases in income will *ceteris paribus* lower the retirement age.

In figure 10, we examine how the uncertainty effect varies with the degree of risk aversion. For each value of the coefficient of relative risk aversion, σ , we choose a value of the utility from leisure, γ , such that the optimal retirement age when life expectancy is 75 is equal to 65. We then consider the effect of changing life expectancy. As the figure shows, the more curved is the utility function (that is, the higher is σ), the less optimal retirement age responds to changes in life expectancy.²⁰

Figure 11 shows the results of combining changes in wages and mortality. We set the coefficient of relative risk aversion to three. The three panels of the figure consider three different levels of the wage: 1930, 1980, and 1990. Within each panel, we look at expected lifetime utility as a function of retirement age for four different sets of life table values: 1900, 1930, 1980, and 1990. Again, we use data from 1980 to calibrate the key parameter γ . Specifically, we choose γ so that, given the 1980 wage and 1980 mortality probabilities, the optimal retirement age in 1980 is 65.

The figure makes it clear that the two effects work in a complementary fashion. For a given level of mortality, an increase in the wage makes retirement more optimal; and similarly, for a given level of the wage, a decline in mortality makes retirement more optimal. Given the 1980 wages, for example, planning for a retirement is not optimal using mortality probabilities from 1930, but it is optimal using mortality probabilities from 1990; similarly, given 1980 mortality probabilities, retirement is not optimal given the 1930 wage, but is optimal given the 1990 wage.

Experimenting with other values of the coefficient of relative risk aversion leads to the conclusion that as risk aversion rises, the income effect increasingly comes to dominate the uncertainty effect as the major cause of retirement. However, using higher values of the coefficient of relative risk aversion also leads to a somewhat troubling conclusion: the model using these values of the risk aversion parameter generates dramatically falling age at retirement as income rises further. That is, unlike the uncertainty effect (which predicts a one-time rise in retirement, followed by rough constancy of the fraction of life spent retired), the income effect predicts that retirement will come to represent an ever larger fraction of life as income rises.

²⁰The values for the other parameters are, $r = 0.06, \theta = 0.03$.

5 Conclusion

Our paper has shown how a reduction in mortality can lead to a shift in the life-cycle pattern of labor supply. High mortality leads to uncertainty about the age of death, and in this environment individuals will find it optimal to work until they die. As mortality falls, it becomes optimal to plan for a period of leisure at the end of life—that is, for retirement. We show, using data on actual changes in the life table over the last century, that this “uncertainty effect” can more than compensate for the more intuitive effect of higher life expectancy in *raising* the retirement age.

As we stressed in the introduction, we do not think of the uncertainty effect as the only explanation for the rise in retirement that has taken place over the last century. Changes in government policy and productive technology, as well as the effect of higher income in raising the demand for leisure, all play a role. Further, interactions of several of these channels are likely to be important. For example, we have shown that the uncertainty effect interacts with the income effect to produce a larger reduction in labor force participation than either channel separately. A project for future work in this area will be to apportion causality for the increase in retirement between the different causes we have discussed.

A second dimension along which the model can be extended is to examine how changes in labor supply feed back, via increased life cycle savings, into higher levels of income. That is, one could marry our model of endogenous retirement and savings with a growth model. In contrast to the standard Overlapping Generations model, in which the fraction of life spent working is fixed, our model suggests that the emergence of retirement, and thus of life cycle saving, may be one of the key steps in the process of modern economic growth.

References

- Carroll, Christopher D. (1997), "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis," *Quarterly Journal of Economics* 112, 1–55.
- Carter, Susan B., and Richard Sutch (1996), "Myth of the Industrial Scrap Heap: A Revisionist View of turn-of-the-century American Retirement," *Journal of Economic History* 56, 5–38.
- Costa, Dora L. (1998), *The Evolution of Retirement: An American History, 1880-1990*, Chicago, IL: The University of Chicago Press.
- Danziger Sheldon, Robert Haveman and Robert Plotnick (1981), "How Income Transfer Programs Affect Work, Savings, and the Income Distribution: A Critical Review," *Journal of Economic Literature* 19, 975–1028.
- Eckstein, Zvi, Pedro Mira, and Kenneth I. Wolpin (1998), "A Quantitative Analysis of Swedish Fertility Dynamics: 1751-1990," CEPR Discussion Paper 1832.
- Ehrlich, Isaac and Francis T. Lui (1991), "Intergenerational Trade, Longevity, Intrafamily Transfers and Economic Growth," *Journal of Political Economy* 99, 1029–1059.
- Fields, Gary S. and Olivia S. Mitchell (1984), *Retirement, Pensions, and Social Security*, Cambridge: MIT Press.
- Graebner, William (1980), *A History of Retirement: The Meaning and Function of an American Institution, 1885-1978*, New Haven, CO: Yale University Press.
- Gruber, Jonathan and David Wise (1998), "Social Security and Retirement: An International Comparison," *American Economic Review* 88, 158–163.
- Hurd, Michael D. (1990), "Research in the Elderly: Economic Status, Retirement, and Consumption, and Savings," *Journal of Economic Literature* 28, 565–637.
- Haines, Michael R. (1994), "Estimated Life Tables for the United States, 1850-1900," NBER Working Paper Series on Historical Factors in Long Run Growth, Historical Paper 59.
- Kalemli-Ozcan, S. (2001). "A Stochastic Model of Mortality, Fertility and Human Capital Investment," *Journal of Development Economics*, forthcoming.

- Keyfitz, Nathan and Wilhelm Flieger (1990), *World Population Growth and Aging*, Chicago, IL: The University of Chicago Press.
- Kotlikoff, L.J., and L.H. Summers (1988), "The Contribution of Intergenerational Transfers to Total Wealth: A Reply," in: Kessler D. and A. Masson eds., *Modeling the Accumulation and Distribution of Wealth*, New York, NY: Oxford University Press.
- Lee Chulhee (1996), "The Expected Length of Retirement and Life-Cycle Savings, 1850-1990," University of Chicago, Working Paper.
- Lee Chulhee (1998), "Life-Cycle Savings in the United States, 1900-1990" Seoul National University, Working Paper.
- Long, Clarence (1958), *The Labor Force Under Changing Income and Employment* Princeton: Princeton University Press for NBER.
- Lumsdaine, Robin, and David Wise (1994), "Aging and Labor Force Participation: A Review of Trends and Explanations," in: Yukio Noguchi and David A. Wise , eds., *Aging in the United States and Japan: Economic Trends*, Chicago, IL: University of Chicago Press.
- Matthews, R.C.O, C.H. Feinstein and J.C. Odling-Smee (1982), *British Economic Growth, 1856-1973*, Stanford, CA: Standford University Press.
- Meltzer, David (1992), "Mortality Decline, the Demographic Transition and Economic Growth," Ph.D Dissertation, University of Chicago.
- Modigliani, Franco (1986) , "Life Cycle, Individual Thrift, and the Wealth of Nations," *American Economic Review* 76.
- Moen, John R. (1988), "Past and Current Trends in Retirement: American Men from 1860 to 1980," *Federal Reserve Bank of Atlanta Economic Review* 73, 16–27.
- Moen, John R. (1994), "Rural Nonfarm Households: Leaving the farm and the Retirement of older Men, 1860-1980," *Social Science History* 18, 55–75.
- Ransom, Roger and Richard Sutch (1986), "The Labor of Older Americans: Retirement of Men on and off the Job, 1870-1937," *Journal of Economic History* 46, 1–30.

- Ransom, Roger and Richard Sutch (1988), “The Decline of Retirement in the Years before Social Security: US Retirement Patterns, 1870-1940,” in: Rita Ricardo-Campbell and Edward P. Lazear , eds., *Issues in Contemporary Retirement*, Stanford, CA: Stanford University Press.
- Sala-i-Martin, X., (1996), “A Positive Theory of Social Security,” *Journal of Economic Growth* 1, 277–304.
- Weil, David N., (1997), “The Economics of Population Aging,” in: Mark R. Rosenzweig and Oded Stark , eds., *Handbook of Population and Family Economics*, North-Holland: Elsevier Science.

Table 1: Optimal Retirement Age

Parameters	γ	1900	1950	1980	2030
$\sigma = 1$					
$r = 0.04, \theta = 0.02$	0.96	∞	∞	55	55
$r = 0.05, \theta = 0.03$	0.90	∞	∞	54	54
$r = 0.06, \theta = 0.03$	0.67	∞	∞	57	57
$r = 0.07, \theta = 0.05$	0.82	∞	∞	53	53
$r = 0.06, \theta = 0.05$	1.07	∞	∞	49	49
$\sigma = 2$					
$r = 0.04, \theta = 0.02$	1.25	∞	∞	54	54
$r = 0.05, \theta = 0.03$	1.12	∞	∞	53	53
$r = 0.06, \theta = 0.03$	0.75	∞	∞	57	56
$r = 0.07, \theta = 0.05$	0.94	∞	∞	52	52
$r = 0.06, \theta = 0.05$	1.40	∞	∞	46	47
$\sigma = 3$					
$r = 0.04, \theta = 0.02$	1.33	∞	∞	58	58
$r = 0.05, \theta = 0.03$	1.16	∞	∞	57	57
$r = 0.06, \theta = 0.03$	0.77	∞	∞	59	58
$r = 0.07, \theta = 0.05$	0.96	∞	∞	55	54
$r = 0.06, \theta = 0.05$	1.44	∞	∞	51	51

Figure 1




Figure 2




Figure 3
Probability of Death Conditional on Reaching 20




Figure 4
Optimal Retirement Age




Figure 5
Maximum Expected Utility




Figure 6
Optimal Consumption Path [$R^*=57$, Life Table 1980]




Figure 7
Optimal Asset Path [$R^*=57$, Life Table 1980]




Figure 8
Optimal Consumption Path [$R=\infty$, Life Table 1980]




Figure 9
Optimal Asset Path ($R=\infty$, Life Table 1980)




Figure 10
Optimal Retirement with CRRA Utility




Figure 11
Interaction of Income and Uncertainty

