The Costs of Uncertain Life Span

Ryan D. Edwards*

June 12, 2006

Abstract

Much uncertainty surrounds the length of life. Based on period rates, the standard deviation in the adult age at death in the U.S., S_{10}, has hovered around 15 years since 1960, while the average adult life span, $\mathrm{M}_{10} \equiv e_{10} + 10$, has increased steadily from about 72 years to 78. In this paper, I show that for reasonable parameter values, this variance in life span represents a significant cost. I construct a simple model that shows each standard deviation in life span is currently worth about half a life year. Numerical simulations suggest this estimate is conservative, and they demonstrate that even full annuitization cannot offset the costs of life-span variance. Viewed in this light, differences in S_{10} and differences in life expectancy are of comparable importance in assessing well-being. An extension is that procyclical fluctuations in mortality, although undeniably costly to some, translate into extremely small increases in S_{10} and thus barely any real cost.

*PRELIMINARY. PLEASE DO NOT CITE OR QUOTE WITHOUT PERMISSION. Research Associate, UC Berkeley Demography, redwards@demog.berkeley.edu. I thank Shripad Tuljapurkar, Ronald Lee, David Canning, and seminar participants at the Harvard Initiative for Global Health for helpful comments. This research was partially supported by NICHD grant T32 HD 07329.

1 Introduction

That life is uncertain is a truism. Uncertainty in the length of life impacts well-being viscerally, and we are certainly aware of it. None of us knows exactly when we will die, and we are all grimly aware of how sudden and unexpected or how drawn-out the process of dying can be.

As discussed by Edwards and Tuljapurkar (2005), we can quantify uncertainty in life span by interpreting the deaths in a period life table as probability densities. Panel A in Figure 1 displays the density function of life spans for the U.S. in 1999, which appears as a skew-left distribution around a mode of 85 years plus a spike at infant mortality. The unconditional mean of this distribution, also called period life expectancy at birth, or e_0, is about 77, while the mean conditional on survival to any early age past infancy, say age 10, is $\mathrm{M}_{10} = 78$. As is plainly visible in the figure, there is considerable variation around the mean even if we omit infant mortality, which is fixed at age 0. The standard deviation in life span after age 10, S_{10}, is about 15 years, producing a coefficient of variation in adult life span of 0.2.

Edwards and Tuljapurkar show that U.S. levels of S_{10} are higher than those in other industrialized nations by between 1 and 3 years. They also discuss how long-term trends in S_{10} suggest relative stagnation rather than the monotonic improvement we see in infant mortality, e_0, and M_{10}. Industrialized nations appear to be stuck with uncertainty in life span, but there are clear differences between high-variance and low-variance countries. To better understand well-being and behavior, we would like to know the costs associated with S_{10}.

A wealth of research exploring the valuation of life (Rosen, 1988; Viscusi, 1993; Johansson, 2002; Aldy and Viscusi, 2003) has produced many insights regarding the willingness to pay for mortality reduction. Closely related work examines the willingness to pay over historical periods and across geographical boundaries (Viscusi and Aldy, 2003; Costa and Kahn, 2004), and other research gauges the value of long-term increases in life span (Cutler and Richardson, 1997; Nordhaus, 2002; Murphy and Topel, 2005; Becker, Philipson and Soares, 2005). Many of these efforts have modeled survivorship realistically, i.e., with variance around mean life span, but they have not explored the costs of variance as distinct from the benefits of the mean. Sometimes this focus is driven by the data. In their study of world development since 1960, Becker, Philipson and Soares (2005) use only mean life spans to value overall well-being because demographic data is extremely scarce for developing countries. But given that life-span variance is surely higher in developing countries, omitting its costs may produce an inaccurate comparison if those costs are substantial. Similarly, vast declines in S_{10} in the industrialized world prior to 1960 (Edwards and Tuljapurkar, 2005) may affect long-term comparisons such as those of Nordhaus (2002) if variance is costly.

To be sure, it is not trivially clear that variance in life span is necessarily costly. Annuities and life insurance are two examples of market instruments that lessen risks associated with uncertain life spans. In this paper, I show that while full annuitization removes all consumption risk associated with mortality, and therefore improves welfare, annuities do not remove the utility costs of S_{10}. Whether life insurance helps offset S_{10} is more difficult to say, since it affects individual utility only through the bequest motive. If people are completely altruistic and are able to insure against income loss associated with early death, it is conceivable that they do not care about variance in life span, or even the mean age at death, for that matter. But full altruism is not a typical assumption or finding in the value-of-life literature, and evidence suggests that bequests are generally not intended (Hurd, 1987, 1989).

Another motivation to consider the costs of life-span variance derives from a separate literature that studies temporal fluctuations in health and mortality. Ruhm (2000, 2003, 2004), Tapia Granados (2005), and others have shown that mortality rates tend to rise above trend during economic expansions and fall during recessions. While growth in incomes and life spans is correlated over the long run, this research shows that in the short run, rapid growth can actually shorten life spans through a variety of channels. The costs of such procyclical mortality, which appear to be borne by a broad range of subgroups rather than a vulnerable few, are obviously high for those afflicted. But what are the costs of procyclical mortality for the population as a whole? It turns out that we can interpret procyclical mortality as another type of life-span variance. Using the framework developed in this paper that estimates the costs of total variance, we can also estimate the costs of procyclical mortality because it is an isomorphic dynamic.

In the sections that follow, I present several models of intertemporal choice that help us understand the costs of variance in adult life spans. I begin by outlining the individual's lifetime planning problem, and I discuss the key parameters, paying special attention to the distribution of life spans. I then propose and solve a highly simplified analytical model that produces an elegant and parsimonious rule expressing the tradeoff between the mean and the variance in life span. Next, I construct numerical simulations of the complete model with fully endogenous consumption, and I test the realism of the simple model by specifying fully realistic mortality, allowing full annuitization, and testing an array of other parameter values.

2 Intertemporal choice under life-span uncertainty

In order to assess the costs of variability in life spans upon individuals, I adopt a standard framework of intertemporal optimization with time-separable utility. This model, which economists use to address a vast array of issues, including the value of life, turns out to imply a specific tradeoff between mean and variance in life span that is a function of the parameters. The price of a standard deviation in life span in terms of the mean takes a convenient form when wealth is fully annuitized. Later I show that numerical solutions of the model with or without annuities confirm this rule of thumb.

Consider an expected utility maximizer at time $t = 0$ with an implicit rate of time discounting equal to δ and no bequest motive. Lifetime expected utility is the sum of period utilities drawn from consumption, $u(c(t))$, weighted by the force of time discounting, $e^{-\delta t}$, and the probability that the individual is alive, $\ell(t)$:

$$EU = \int_0^\infty u(c(t))e^{-\delta t}\ell(t)\ dt. \tag{1}$$

Demographers call $\ell(t)$ the survivorship function, and it is shown in Panel B of Figure 1. The decrement to $\ell(t)$, which is one minus the cumulative density function of life span, is the probability density function shown in Panel A of Figure 1 and commonly referred to as life-table deaths, $d(t)$. It is also the probability distribution of life span. Finally, we define the life-table probability of dying between t and $t+1$ as $q(t) = \log \ell(t) - \log \ell(t+1)$.[1] Mortality tends to increase exponentially with age (Gompertz, 1825), as shown in Panel C of Figure 1.

Suppose the individual has a financial endowment W that can be consumed or saved at a fixed market rate of interest, r. For simplicity, there is no labor, capital, financial risk, or human capital in this model.[2] In a

[1] This form is more convenient for expositional purposes here, but life-table $q(t)$ is typically defined implicitly as $\ell(t+1) = \ell(t)[1 - q(t)]$, in order to attrit the entire cohort at a finite age. In continuous time, $q(t)$ is the hazard or mortality rate.

[2] Kalemli-Ozcan, Ryder and Weil (2000) and Li and Tuljapurkar (2004) develop models that include retirement, endogenous capital accumulation, and education alongside mortality. Each element probably increases the costs of life-span variance. If individuals must trade their leisure time for market earnings, higher S_{10} erodes expected lifetime wealth provided that the retirement age is within the support of probabilistic life span. If capital

market without annuities, the budget constraint requires the individual to finance the present value of all future consumption out of wealth:

$$W = \int_0^\infty c(t)e^{-rt}\ dt. \tag{2}$$

With this budget constraint, the model produces unintended bequests of varying size for all individuals except those who live forever, which occurs with probability zero. Bequests are larger the earlier death occurs. If instead actuarially fair annuities are available, the budget constraint takes on a different form (Becker, Philipson and Soares, 2005):

$$W = \int_0^\infty c(t)e^{-rt}\ell(t)\ dt. \tag{3}$$

An annuity pays off in future periods only if the individual is alive. This allows the buyer to finance future consumption more cheaply, at the expense of unintended bequests.[3] Although public annuities like Social Security exist, private annuities markets are notoriously thin (CITE), so partial annuitization best describes reality. But the extreme cases are easier to model, and it will turn out that the presence or absence of annuities does not affect the costliness of life-span variance very much. I discuss the impact of annuities on the costs of life-span variance in greater detail below.

The individual maximizes equation (1) subject either to (2) or (3) depending on whether annuities are available. I can write the resulting Euler condition that describes intertemporal choice as

$$u'(c(t+1)) \quad = \quad u'(c(t))e^{\delta - r + D\cdot q(t)}, \tag{4}$$

where D is an indicator of the lack of annuities. With full annuitization, $D = 0$ and mortality cancels out of (4) because it appears in both the objective and the constraint under full annuitization. When $u(\cdot)$ takes the familiar form of power utility with constant relative risk aversion over consumption,

$$u(c(t)) = \frac{c^{1-\gamma}}{1-\gamma} + K, \tag{5}$$

and the interest rate were endogenous, higher S_{10} would likely deplete the capital stock by lowering the marginal utility of wealth, raising the interest rate and lowering the wage rate. Effects on welfare are countervailing, but it seems likely that the net effect would be negative. Human capital investment is riskier when S_{10} is higher, which should result in lower educational attainment and a decrease in welfare.

[3] In this model with costlessly enforced contracts, the price of the annuity is the right to leave bequests. All wealth that is unused by those who die is redistributed to the living.

for some constant K, equation (4) implies

$$c(t) \quad = \quad c(0)e^{t[r-\delta-D\cdot q(t)]/\gamma} \tag{6}$$

where $c(0)$ is a function of wealth, the parameters, the moments of life span, and the annuitization indicator, D. Without annuities, the exponentially increasing mortality rate, $q(t)$, ultimately pulls consumption lower through age, producing a consumption trajectory that looks like the survivorship curve. Annuitization removes this effect. When $r = \delta$, consumption under annuitization is flat through age. The utility curvature parameter γ magnifies or dampens the impacts of differences in r and δ and of mortality under no annuities.

Under full annuitization, we can completely solve the model by reformulating the budget constraint:

$$W \quad = \quad E_T\left[\int_0^T c(t)e^{-rt}\ dt\right], \tag{7}$$

where T is the realization of life span, and we have removed the survivorship weights and instead taken the expectation over T. For a given distribution of T, we can use (6) with $D = 0$ to solve the integral within the expectation and then find $c(0)$ by taking the expectation. Without annuities, the presence of $q(t)$ in the consumption function precludes analytical solutions because mortality increases exponentially with age. For now, I focus on the case of full annuitization. Later, I use numerical simulation to examine the effect of annuitization.

We are interested in the relative price of variance in life-span in this model, which is governed by its marginal utility relative to that of others. To proceed, we can rewrite lifetime expected utility as

$$EU = E_T\left[\int_0^T u(c(t))e^{-\delta t}\ dt\right]. \tag{8}$$

Using the power utility formulation in (5) and the consumption function in (6) with $D = 0$, we find that lifetime utility under full annuitization is

$$EU \quad = \quad E_T\left[\frac{c(0)^{1-\gamma}}{(1-\gamma)\hat{\delta}}\left(1-e^{-\hat{\delta}T}\right)+\frac{K}{\delta}\left(1-e^{-\delta T}\right)\right], \tag{9}$$

where

$$\hat{\delta} = \delta - \frac{1-\gamma}{\gamma}(r-\delta). \tag{10}$$

When r is close to δ, we have $\hat{\delta} \approx \delta$, with equality when $\gamma = 1$. Equation (9) implies that individuals are risk averse over life span, T, in this model. The Arrow-Pratt coefficient of absolute risk aversion over T, $-EU_{TT}/EU_T$, is approximately equal to the rate of time discounting, δ, and exactly equal when $r = \delta$. Absolute risk aversion in life span is roughly constant.[4]

We would expect individuals who are risk averse over life span to be hurt by uncertainty in life span, and this is exactly what we see if $T \sim N(M, S^2)$.[5] When life spans are normal, lifetime utility is

$$EU = \frac{c(0)^{1-\gamma}}{(1-\gamma)\hat{\delta}} \left[1 - e^{-\hat{\delta}M + \hat{\delta}^2 S^2/2}\right] + \frac{K}{\delta}\left[1 - e^{-\delta M + \delta^2 S^2/2}\right], \quad (11)$$

Expected lifetime utility is a decreasing function of S, provided that period utility is positive.[6]

We can find the price p_S of a standard deviation in life span, S, in terms of the mean, M, by constructing the ratio of their marginal lifetime utilities. When we consider the case of $K = 0$, a parsimonious formula emerges:[7]

$$p_S = \frac{\partial EU/\partial S}{\partial EU/\partial M} = \frac{-\hat{\delta}^2 S \; EU}{\hat{\delta} \; EU} = -\hat{\delta} S. \quad (12)$$

This formula is exact for infinitesimal changes in the moments of life span under annuitization. The price is negative when $\hat{\delta} > 0$ because S is a bad. In order to be just as happy with higher variance, the individual requires a

[4] Constant absolute risk aversion over life span seems consistent with time trends in life expectancy and in S_{10} duing the modern era, just as constant relative risk aversion over consumption risk matches the patterns in aggregate consumption. Average life spans have grown linearly over time, with a fairly constant standard deviation in levels, while consumption grows exponentially with steady percentage deviations.

[5] As shown by Figure 1, adult life spans are not technically normal, with both leftward skewness and leptokurtosis, or peakedness with fat tails, an indicator of different subgroup variances. Numerical simulations show that normality reduces the costs of S.

[6] It is a standard observation that period utility, which is the marginal utility of being alive in that period, should be nonnegative when modeling dynamics of life span (Rosen, 1988; Hall and Jones, 2004; Becker, Philipson and Soares, 2005). If it were negative or zero, a utility maximizing individual would choose to die. Becker, Philipson and Soares (2005) calibrate the additive utility shifter $K < 0$. This does reduce technically the cost of S through the second piece of (11), but numerical simulations confirm this effect to be small and uninteresting.

[7] The constant utility shifter does not appreciably augment the insights to be gained. On its own, it implies the same dynamics as when $r = \delta$, with only δ mattering for costs. When combined with the flow utility from consumption, both numerator and denominator in (12) are weighted averages of the two pieces in equation (11). When the piece with K has more weight, the coefficient on S in p_S is weighted more toward δ than $\hat{\delta}$.

higher mean. We also see that the magnitude of p_S increases linearly with the level of S, with slope equal to $-\hat{\delta}$. As shown by equation (10), the magnitude of the slope, $\hat{\delta}$, is an increasing function of δ if $\gamma > 0$; it is a decreasing function of r when $0 < \gamma < 1$ and an increasing function of r when $\gamma > 1$; and it is a decreasing function of γ if $\delta > r$ and an increasing function of γ if $r > \delta$.

The implications of equation (12) for population health are significant. In 1999, individuals in the U.S. experienced a standard deviation in life spans conditional on survival to age 10, S_{10}, equal to 15 years. At that level, each year of standard deviation is worth about 0.45 year of mean life span in this model, assuming $r = \delta = 0.03$, its standard value in calibration exercises (Hubbard, Skinner and Zeldes, 1994; Becker, Philipson and Soares, 2005). In Sweden that same year, S_{10} was about 13. If $p_S = -0.45$, individuals in the U.S. would be willing to give up almost 0.9 year in mean life span to have the lower S_{10} of their Swedish counterparts.[8]

Although 0.9 year sounds small relative to mean life span in either country, it is large when compared to the difference in means between the two countries. The mean life span conditional on survival to age 10, M_{10}, was 77.7 years in the U.S. and 80.0 in Sweden in 1999. After adjusting for differences in S_{10}, the total difference in population health between the U.S. and Sweden was more like 3.2 years per person rather than 2.3, an increase of more than a third.

2.1 Numerical solutions of the full model

With numerical methods, we can examine how p_S changes under alternative assumptions about wealth annuitization and mortality. In this section, I model several alternative scenarios in greater detail using numerical methods, with parameters initially set to match those used by Becker, Philipson and Soares (2005): $r = \delta = 0.03$, $\gamma = 0.8$, and $K = -16.2$. I set initial wealth at \$800,000, which is consistent with $r = 0.03$ and per capita consumption of \$26,650 per year, and I fully endogenize consumption. I simulate the model with full annuitization and then the model with no annuitization and compare both sets of results to those implied by the analytical model with annuitization.

For better tractability and for clearer comparisons, I begin by modeling life span as normally distributed. I set mean and variance equal to U.S.

[8] For large changes in the moments, S and thus p_S will change. The isoquants of lifetime expected utility for two normally distributed life spans $L_1 \sim N(M_1, S_1^2)$ and $L_2 \sim N(M_2, S_2^2)$ are given by $M_1 - M_2 = \hat{\delta}(S_1^2 - S_2^2)/2$.

levels in 1994, $M = 76.85$ and $S = 15.66$ years, and I search for the mean life span that compensates expected utility for a decrease in S to 15.05, the 1999 level. Later, when I implement realistic mortality, it is convenient to use life-span distributions from these two years, because there is a relatively large difference in S but a small difference in M. I truncate these synthetic distributions at ages 0 and 150 and rescale so that their cdf's sum to unity.[9]

2.1.1 Normally distributed life spans

Figure 2 plots p_S as given by equation (12) on the same axes with two other loci that I obtain from numerical simulation of the model with normally distributed life spans. I fixed the rate of market interest at $r = 0.03$ and examine how varying δ, which is shown along the horizontal axis, changes p_S. The thick solid line in Figure 2 is $p_S = -\hat{\delta}S$ from the analytical model. The steeper, more curved thin solid line shows the locus obtained from the numerical model with full annuitization of wealth, while the lower dashed line depicts the schedule that results from the numerical model without annuitization.

The two solid lines depict some differences between the analytical and numerical versions of the model with annuities. They cross at $\delta = 0.03$, and the numerical model produces a more steeply sloped p_S locus that appears to be concave at high δ. It turns out that these differences are the result of the utility shifter K in the numerical model, which when negative tends to increase $\hat{\delta}$ and thus p_S when δ is large. When $K > 0$, the numerical locus is flatter than the analytical locus, and when $K = 0$, the lines overlap.

Figure 2 reveals a more interesting relationship between the costs of uncertainty under annuitization and the costs without annuitization. The dashed line, which shows p_S without annuities, is significantly below the other two, indicating that life-span variance is more costly at any δ. At the baseline of $r = \delta = 0.03$, p_S without annuitization is about -0.75, more than half again as large as p_S with annuitization, which is about -0.45. This is an intuitive result insofar as annuitization is designed to remove risks to consumption associated with uncertain life span. That annuitization

[9] Age 150 is an unrealistic but convenient choice when life spans are normally distributed. The Human Mortality Database (2006) topcodes age at 110, and there are few documented individuals who have survived to that age. When life spans are normally distributed with means around age 80, densities past age 110 are not miniscule. Truncating at age 110 actually creates significant skewness in the distribution, and probably changes the mean and variance. Such skewed distributions actually produce a p_S locus that better resembles that produced by realistic survivorship. This is likely because life spans are skew-left in reality.

removes only about one third of the total costs of life-span uncertainty under baseline parameter values is more surprising. In this model with no bequests, the direct utility costs of S are more important than the costs of consumption uncertainty as long as δ is not considerably less than r.

We can also vary r while holding δ fixed, which is shown in Figure 3. three p_S schedules, this time as functions of r with $\delta = 0.03$. Here, an even greater difference emerges between the numerical and analytical models under annuitization, which is again linked to the utility shifter K. The numerical model reveals considerably more negative p_S when $r < \delta$ and $K < 0$. When $r < \delta$, consumption optimally declines over time. If in addition $K < 0$, the benefit of living long is weakened considerably relative to the cost of dying early, so the costs of S are higher. As before, we see that p_S is more negative when annuities are not available, regardless of the rate of market interest.

Each locus shows that the costs of S are declining in magnitude with r. If $0 < \gamma < 1$, a very large r can actually produce a negative $\hat{\delta}$, around $r = 0.15$ when $\gamma = 0.8$ and $\delta = 0.03$. When r is high, the price of future consumption is low. If intertemporal substitution is elastic, when $0 < \gamma < 1$, the individual plans more future consumption. Since the marginal utility of a year of life is equal to the period utility function, this increase in future consumption raises the marginal utility of mean life span and lowers p_S.

2.1.2 Realistic life spans

Modern distributions of human life span are skew-left and leptokurtic. The skewness is of direct interest here, because it implies that a mean-preserving spread in life span lowers survivorship probabilities asymmetrically. Colloquially, we could describe an increase in variance when there is leftward skewness as reducing survivorship at young adult and adult ages, while increasing it only at old ages. Since an individual with positive time discounting values the present more than the future, we expect skewness to result in higher true costs of uncertain life spans.

It is tricky to model realistic life spans with particular means and variances because we do not have a convenient functional form of the probability distribution. I proceed by generating additive translations of the 1999 life-span distribution above age 10 in the U.S., which originally had a mean of 77.67 and a standard deviation of 15.05, so that I have an array of realistic distributions with varying means but fixed variances. Then I search for the distribution that produces the same lifetime expected utility as the 1994 distribution above age 10 with $\mathrm{M}_{10} = 76.85$ and $\mathrm{S}_{10} = 15.66$.[10]

Figure 4 is the realistic-mortality analogue of Figure 2, depicting three loci of p_S against δ for $r = 0.03$, using the same vertical scale for easier comparison. The thick solid line, which shows the analytical model's p_S, is the same as before. The other two schedules reveal levels of p_S that are more negative than in Figure 2, especially for larger δ. The costs of life-span uncertainty are indeed higher when I model life span realistically, with leftward skewness that places a wider range of early years at risk. The thin solid line, representing the output of the numerical model with annuities, shows $p_S = -0.7$ when $\delta = 0.03$, greater in magnitude than the $p_S = -0.45$ from the simple model. As δ rises, both numerical models show precipitous increases in the costs of variance.

Some human life spans end at birth or shortly thereafter. The variance in life span attributable to infant mortality is quite costly, but I have not focused explicitly on its costs here. This is because modeling infant mortality is in a sense complex and also easy. We do not have a model for how improvements in adult health, which clearly lengthen adult life spans and may also decrease variance, may also lower infant mortality. At the same time, it is clear that a reduction in infant mortality, if we treat it as entirely separate from adult mortality, has more obvious benefits than a reduction in adult life-span variance. One less death is one more average adult life span, other things equal. No more finesse is needed.

Still, the presence of infant mortality may affect p_S, since infant mortality reduces the length and raises the total variability of life. By making mean life scarcer, infant mortality may lower p_S because the marginal utility of M is higher. But we have seen that the marginal utility of S rises with S, so a rise in infant mortality and total variance may have the same effect, which would tend to increase p_S.

To assess the costs of adult variance under fully realistic survivorship with infant mortality, I simply treat infant mortality as a completely separate dynamic. I fix infant deaths at their relative probability in 1999 and reestimate the compensating change in mean adult life span that offsets the reduction in S from 15.66 to 15.05.[11]

[10] Using MATLAB, I apply a cubic spline to the distribution of life spans by single years of age over age 10 in the U.S. in 1999, and I sequentially evaluate the spline at hundredths of a year in age, spaced one year apart. I then redefine age back to whole years, which produces a sideways translation of the distribution, changing the mean but retaining the same variance. At ages under 10, I simply duplicate the density at 10 and renormalize the entire distribution. Later I include realistic infant mortality as discussed in the text.

[11] I overlay the life span distribution under age 10 in 1999 on top of the distribution in 1994 and on top of each translated distribution from 1999 that has a different mean.

Figure 5 depicts p_S as functions of δ when $r = 0.03$. Survivorship is now fully realistic, with infant mortality in addition to leftward skewness and leptokurtosis in the adult hump. Including infant mortality does not significantly change the qualitative results, but in the numerical models p_S is generally smaller in magnitude than it was in Figure 4 with realistic adult mortality but no infant mortality. Including infant mortality increases the marginal utility of mean life and not the marginal disutility of a standard deviation in life span, so p_S is less negative than before.

3 How uncertain are actual life spans?

We have seen that the price of a standard deviation in life span in terms of years of life is roughly the current level of standard deviation times the rate of time discounting, δ. Measured in terms of standard deviation, the more uncertainty there is, the higher its cost. Since the level is clearly important, it is worth discussing whether the measure I have used so far, S_{10}, the standard deviation in life span above age 10 based on period rates, is indeed the correct measure for individuals and populations.

3.1 Periods, cohorts, and forecasts

Up to now, I have proxied the actual, or cohort life spans of individuals with those based on period mortality rates. The period standard deviation in life span, S_{10}, measures uncertainty for a synthetic cohort of individuals living their entire lives in the relevant period. We frequently measure population health using period life expectancy at birth, e_0, which is the average life span based on period rates and also applicable to a synthetic cohort. Period S_{10} is in this sense the appropriate second moment of health for the adult population. The costs of period S_{10} are felt by an average individual alive today in an age-adjusted sense, just as the benefits of an increase in period life expectancy reflect an average benefit for living individuals.

To be sure, there are clear conceptual differences between cohort and period standard deviations in life span, but differences are not overwhelmingly large in an ex post sense. Figure 6 plots both period and cohort S_{10} by calendar year or birth year for the U.S. since 1900. I compute period S_{10} using annual period life tables from the Human Mortality Database (2006), while I form cohort S_{10} using decennial cohort life tables collected and forecast by the Social Security Administration (Bell and Miller, 2005). Data

Then I renormalize so that each cdf sums to unity.

for only the very oldest birth cohorts shown, those who are now completely extinct, are real. For cohorts with any surviving members, S_{10} is at least partially a forecast based on future trends projected by the SSA. In the figure, both period and cohort S_{10} drop from around 20 years to around 15 years over the last 100 years and remain there. As described by Edwards and Tuljapurkar (2005), period S_{10} fell during the epidemiological transition that ended around 1960, as infectious diseases decreased in relative risk. We see the very same pattern in cohort S_{10} plotted by birth year presumably because infectious disease raised S_{10} the most through exposure at younger ages.

But this view misses a subtler source of uncertainty in life span for the individual. We typically use period rates to assess well-being because we do not know with certainty the age-specific mortality rates that will apply to living cohorts over their lives. Lee and Carter (1992) model age-specific mortality in the U.S. over the last century, and they find that rate of decline in mortality is a random walk with drift. In addition to the uncertainty in life span stemming from nonzero mortality rates at every age in a known life table, which we can label as life-table uncertainty, an individual also faces considerable forecast uncertainty. Over an average life span, forecast uncertainty can become quite large. These two sources of variance in life span are completely independent and thus cumulate.[12]

(***Figure/description TBD. It would be nice to quantify Lee-Carter forecast uncertainty for either a period or a cohort or both.)

A living cohort may therefore face life span uncertainty that differs from cohort S_{10} implied by a projected life table. But even with forecast uncertainty, a cohort at some average age may experience roughly its forecast cohort S_{10}, and the average individual alive today will experience roughly the period S_{10}. This is because both life-table uncertainty and forecast uncertainty diminish with age and lower the total uncertainty faced by older individuals. A newborn in the U.S. faces the maximum amounts of life-table and forecast uncertainty, while someone who has survived to age 50 and expects only another 30 years of life faces considerably less of each.

[12]There are two caveats. First, a standard deviation in life span rises only with the square root of the total variance, or less than additively. Second, the forecast may actually predict a change in life-table uncertainty. Neither Lee and Carter (1992) nor Social Security forecasts predict any change in period S_{10}, however.

3.2 Subgroup differences in variance

A separate issue that arises in measuring uncertainty in life span for specific individuals is the fact there is much subgroup heterogeneity in S_{10}. Edwards and Tuljapurkar (2005) show that S_{10} is lower among females than males, and that it is higher among African Americans, individuals with lower household income, and individuals with less education. In many cases these subgroup differences amount to between 1 and 3 years in standard deviation, which implies a difference in p_S of almost 0.1 year if $\delta = 0.03$, an increase of some 20 percent from the full-annuities baseline of $p_S = -0.45$.

Costs of life-span uncertainty are also higher if δ is higher. There is evidence of higher time discounting among individuals with lower socioeconomic status (SES), who save less and tend to engage in behavior that outwardly appears myopic (CITE).[13] Low-SES groups with high S_{10} and higher δ must also bear higher costs of life-span uncertainty.

The less fortunate also bear a heavier burden if they have disproportionately less access to annuities, which we have found to offset perhaps one third of the cost of life-span uncertainty. This characteristic may be observationally linked to high δ, since one reason why low-SES individuals might appear to have high δ, myopia, or insufficient saving, is if they face liquidity constraints or incomplete markets. We would expect that access to annuities markets would also be poor for liquidity constrained individuals.

4 Extension: the costs of procyclical mortality

Recently, a new literature has associated temporarily higher mortality with temporarily better economic times, a pattern of procyclical mortality (Ruhm, 2000, 2003, 2004; Tapia Granados, 2005). The long-term trend in income per capita is exponential growth, but there are business cycle fluctuations

[13] In one sense, this pairing seems counterintuitive. We would expect that those with higher δ and thus higher costs of life-span variance should have chosen less variance, not more. But low-SES individuals have less wealth, which lowers purchases of all goods including less variability. People endowed with low SES also appear to have been endowed with shorter mean life spans, which raises the marginal utility of the mean relative to the standard deviation and results in more standard deviation.

It is far from clear how much of health or life span is determined by SES versus how much of SES is determined by health. Recent work in this field emphasizes the centrality of education and of early life conditions, and at least one switch in the direction of health-SES causality during the life course (CITE) (Smith, 1999). And while we tend to assume that preference parameters like δ are exogenous, we also know that learning patience and self-sacrifice is an important component of human development (CITE). So it is easy to imagine causality running every which way in the nexus of life span, SES, and preferences.

around the trend. Similarly, the secular trend in mortality rates is one of exponential decline with periodic deviations (Lee and Carter, 1992).[14] Procyclical mortality is the temporary departure from trend that researchers have found to be associated with the economic business cycle.

We can assess the economic costs of procyclical mortality using our model of the costs of life span uncertainty.[15] This is because procyclical mortality produces additional uncertainty around life span that is roughly isomorphic to the life-table uncertainty that we have already considered.

Researchers in this subfield find a linear relationship between the change in log mortality for group i at t, $m_{i,t}$, and either the change in log GDP or in some other measure of macroeconomic activity, denoted by V, when ΔV_t is stationary:

$$\Delta \log m_{i,t} = a + b\Delta V_t + \epsilon_{i,t}, \tag{13}$$
$$\Delta V_t = g + \nu_t, \tag{14}$$

where a, b, and g are constants, $\epsilon_{i,t} \sim N(0, \sigma_\epsilon^2)$ and $\nu_t \sim N(0, \sigma_\nu^2)$. Similar to the Lee and Carter (1992) model, equation (13) specifies exponential decline in mortality rates when $a + bg < 0$. There are two sources of temporary fluctuations around the downward trend given by a: $\epsilon_{i,t}$ itself, and procyclical mortality when $b > 0$ and ν_t is transmitted to $\log m_{i,t}$.

As discussed by Edwards (2005), a standard deviation in the annual percentage decline in mortality in the U.S. has been 0.018 since 1947, while a standard deviation in GDP growth has been 0.025. With $b = 0.25$, which is the average coefficient recovered by Tapia Granados (2005), procyclical

[14] Income per capita and mortality rates display different time-series properties, but we can loosely describe both as composed of short-term fluctuations around long-term trends. The formal differences in statistical models is not important for the present exercise. Growth in income per capita is well approximated by an autoregressive model, while the exponential decline in mortality rates is a random walk with drift (Lee and Carter, 1992). Shocks are temporary in the former and permanent in the latter, but the drift term produces behavior that is qualitatively similar to a long-run trend.

[15] One might be tempted to use a simpler representative-agent framework, with a single mortality rate for the whole population, to assess the costs of procyclical mortality. Lucas (1987) assesses the costs of variance in consumption using such a model, where expected utility is $\int e^{-\delta t} u(c(t))\, dt$ and consumption is lognormal. If we included mortality in the standard way in that model, expected utility is $\int e^{-(\delta+m)t} u(c(t))\, dt$, where m is the population mortality rate. A mean-preserving spread in m actually raises expected utility at each t in this model through Jensen's Inequality, since e^{-m} is convex. Variance in mortality is a good in that simple framework. We will see, however, that when mortality rates rise realistically with age, a proportional shock to mortality at every age effectively adds variance in period life spans. Thus procyclical mortality is indeed costly.

mortality accounts for about $b \cdot \sigma_\nu = 0.00625$, or about one third of the annual standard deviation in mortality decline. The rest is absorbed by $\epsilon_{i,t}$.

Some variation in b across ages and sexes emerges in the literature. But to a first approximation, the incidence of procyclical mortality is remarkably homogeneous across groups (Tapia Granados, 2005; Edwards, 2006). To a first approximation, we can then express mortality for age x in a particular period t as a Gompertz (1825) curve with age-independent shocks ϵ_t and $\pi_t = b \cdot \nu_t \sim N(0, b^2\sigma_\nu^2)$, the procyclical mortality shock:

$$\log m_{x,t} = \alpha_t + \beta_t x + \pi_t + \epsilon_t, \tag{15}$$

where both α_t and β_t can vary over time.

Consider the impacts of procyclical mortality π_t on period moments of life span. In equation (15), the procyclical mortality term effectively adds stochasticity to the constant term, α_t, and not the slope, β_t. Tuljapurkar and Edwards (2005) discuss how changes in the Gompertz α_t affect only mean life span while changes in β_t shift the variance.[16] Thus to a first approximation, we can intepret procyclical mortality as imparting stochasticity to mean life span. This procyclical uncertainty around the mean of life span is independent from the life-table uncertainty in life span that produces S_{10}. To rough approximation, these two sources of variance are additive.

To recap, I argue that a standard deviation in the procyclical mortality term π_t, which is about 0.00625, translates into an incremental addition to life-span uncertainty. The next step is figuring out how big an increment it represents.

Vaupel and Romo (2003) derive a relationship between percentage changes in mortality rates and linear increases in mean life span. They show that $\Delta e_0 \approx \bar{\rho} \cdot e^\dagger$, where $\bar{\rho}$ is the average rate of decline in mortality rates, and $e^\dagger$ is the average number of life years lost to a death. Intuitively, this relationship simply asserts that the gain in mean life years is equal to the product of the improvement in mortality and the life years gained per improvement. In the U.S. today, $e^\dagger \approx 12$, since most deaths are distributed around the old-age hump.

Using this formula, we can translate $\sigma_\pi = 0.0065$ into a standard deviation in life span, $S_\pi = 12 \cdot \sigma_\pi = 0.078$. The life-table uncertainty represented

[16] We can see this intuitively by redefining age as $x = \bar{x} + \Delta x$. This amounts to an additive translation of the life-span distribution, which should affect only mean life span and not the variance. In the Gompertz model, redefining age in this way simply adds $\beta_1 \Delta x$ to α_t. Taken together, changes in the Gompertz constant are equivalent to changes in mean life span.

by S_{10} and the procyclical uncertainty in S_π are independent, so the variances are strictly additive. This implies that the increment to S_{10} associated with procyclical mortality is considerably less than S_π and is given by $\left(S_{10}^2 + S_\pi^2\right)^{1/2} - S_{10}$. This increment is extremely small, about 0.0002 when $S_\pi = 0.078$ and $S_{10} = 15.05$ as in the U.S. in 1999.[17]

Viewed in this light, the costs of procyclical mortality are also extremely small. Using the simple relationship derived earlier, $p_S = -\hat{\delta} S$, we can translate an addition to the standard deviation in life span of 0.0002 into mean years of life. When $r = \delta = 0.03$ and $S = 15.05$, this formula suggests that the cost of procyclical mortality is only 0.00009 year of life. At the standard rate of \$100,000 per life year (Tolley, Kenkel and Fabian, 1994), procyclical mortality only costs about \$9 per capita in this framework.

5 Conclusion

Although procyclical mortality may be relatively costless, I have demonstrated that overall uncertainty in life spans is anything but. Life span is a good, and like any other type of good, its value diminishes under uncertainty.

In this paper, I have shown that in a standard model of time-separable utility with no bequest motive, uncertainty in life span is costly when the force of time discounting in utility, δ, is positive When wealth is annuitized and life spans are normally distributed, the price of a standard deviation in adult life span, S, in terms of mean life span, M, is approximately the product of the current level of S and δ: $p_S = -\delta S$, provided that the market rate of interest, r, is near δ. When $\delta = 0.03$, the current level of $S = 15$ in the U.S. implies that we would be willing to give up 0.45 life year in return for one year less in standard deviation.

Numerical simulations confirm that this simple rule remains relatively accurate in more realistic settings when wealth is not annuitized and when life spans are modeled more realistically. When δ is small relative to the market rate of interest, r, costs of S are also small. I also model wealth annuitization and its effects on the costs of life-span uncertainty. At baseline,

[17] I checked the accuracy of these calculations with a numerical simulation using age-specific mortality rates for both sexes combined in the U.S. in 1999. I drew a thousand stochastic shocks s with standard deviation 0.01 for each age in the full distribution, and I multiplied each mortality rate by e^s. Then I calculated a thousand period life tables for each schedule, calculated moments for each, and found that over these thousand distributions, the new standard deviation over age 10 averaged to approximately $\left[S_{10}^2 + (0.01 \times 12)^2\right]^{1/2}$, where S_{10} was the original measure in the period life table.

annuitization of wealth, which eliminates consumption risk due to uncertain life span, removes only about one third of the total costs of life-span uncertainty. The other two thirds represent direct utility costs.

Fully realistic mortality does not significantly alter the results, and it definitely does not diminish the costs. Relative to the convenient but unrealistic case of normally distributed life spans, the leftward skewness that we see in real life spans inflates costs by trading away more valuable young-adult and adult years in return for years in old age. The presence of infant mortality raises the marginal utility of mean years of life, but not by enough to lower the price of a standard deviation in life span in terms of mean years.

I do not account for education or for capital, both of which are considered in a separate study by Li and Tuljapurkar (2004). I also do not consider the costs of variation in morbidity, or the quality of life, and it is easy to imagine that the total cost of variation in both quantity and quality of life may differ from the cost of variation in quantity alone. Risk is compounded when realizations are positively correlated, reduced when they are negatively correlated, and if there is zero correlation, risks are independent. With morbidity and life span, the sign of the correlation is not altogether clear. On the one hand, morbidity frequently precedes death; in this sense the two are positively correlated, implying that their risks cumulate. But sometimes individuals enter into extended periods of morbidity without imminent death, for example in nursing homes. Morbidity and death may be negatively correlated for certain lingering diseases like Alzheimer's or nonlethal but disabling physical trauma. Without more knowledge of the covariance structure, it is difficult to say how variation in the quality of life may affect the costs of life-span variance.

My assumptions about bequests play an important role. Bequests are unintended in the model that I consider without annuities, and they are zero in the model with full annuitization. If there were a bequest motive, it could significantly reduce the cost of life-span uncertainty by compensating early death with the utility of leaving an inheritance. But Hurd (1987, 1989) has shown that bequests are not intentional. Even if they were, it is hard to imagine an individual could fully insure against life-span risk through bequests. Such an individual would not find value in living as opposed to dying beyond procreation, and such behavior is nowhere to be seen in the literature on the value of life.

The practical implications of my findings are several. Variance is an important moment in life-span distributions. The oft-cited difference in mean life spans between the U.S. and Sweden, which was 2.3 years in 1999, is only part of the full difference in population health between the two countries. In addition to a longer average life span, Sweden also enjoys a lower standard deviation in adult life span, $S_{10} = 13$, compared to $S_{10} = 15$ in the U.S. Priced at 0.45 year per year in standard deviation, this difference in the second moment is worth about 0.9 year, or more than a third of the difference in mean life span. Priced at $100,000 per life year, the total difference in life span between the two countries is more like $320,000 per person rather than $230,000.

Another implication is that the benefits of rapid gains in period life expectancy among developing countries during the past 50 years may have been tempered by less progress against variance in life span. Valuing gains in life expectancy in addition to gains in income per capita, Becker, Philipson and Soares (2005) reveal considerable growth in total well-being among developing countries and thus much world-wide convergence. But owing to data constraints, they can only value the changes in the first moment of life span. The extent of life-span variance in developing countries remains unknown, but it is surely higher than in industrialized countries. High-quality data in the Human Mortality Database (2006) show that S_{10} is currently between 16 and 18 in three former Soviet republics including Russia. Although the basic convergence result of Becker, Philipson and Soares probably remains true, it seems likely that poor countries are still considerably poorer than rich countries in part because they have higher variance in life spans.

A third implication is closely related. Nordhaus (2002) and others have valued the progress against mortality in the U.S. over time. Researchers frequently conduct such exercises in order to ascertain whether medical care or research has been worth its financial costs. These studies typically value only changes in the crude death rate or in life expectancy, which according to my results will understate the true benefits if variance in life span is also declining. Since 1960, S_{10} has not fallen much in the U.S., as shown by Figure 6. But there were large decreases in variance during the earlier epidemiological transition (Edwards and Tuljapurkar, 2005).

Finally, it is worth considering the implications of these findings for individuals and for policy. Subgroups with low SES suffer both lower mean life span and higher variance. Their well-being would be significantly improved by more income, more life span, and less variance, since their marginal utilities of each are high.

References

Aldy, Joseph E. and W. Kip Viscusi. 2003. "Age Variations in Workers' Value of Statistical Life." *NBER Working Paper* 10199.

Becker, Gary S., Tomas J. Philipson and Rodrigo R. Soares. 2005. "The Quantity and Quality of Life and the Evolution of World Inequality." *American Economic Review* 95(1):277–291.

Bell, Felicitie C. and Michael L. Miller. 2005. "Life Tables for the United States Social Security Area, 1900-2100." Actuarial Study No. 120, Office of the Chief Actuary, Social Security Administration.

Costa, Dora L. and Matthew E. Kahn. 2004. "Changes in the Value of Life, 1940–1980." *Journal of Risk and Uncertainty* 29(2):159–180.

Cutler, David M. and Elizabeth Richardson. 1997. "Measuring the Health of the U.S. Population." *Brookings Papers on Economic Activity, Microeconomics* pp. 217–271.

Edwards, Ryan D. 2005. "Commentary: Work, Well-Being, and a New Calling for Countercyclical Policy." *International Journal of Epidemiology* 10.1093/ije/dyi142.

Edwards, Ryan D. 2006. "Who Is Hurt by Procyclical Mortality?" Unpublished manuscript available at http://www.demog.berkeley.edu/˜redwards, April.

Edwards, Ryan D. and Shripad Tuljapurkar. 2005. "Inequality in Life Spans and a New Perspective on Mortality Convergence Across Industrialized Countries." *Population and Development Review* 31(4):645–675.

Gompertz, Benjamin. 1825. "On the nature of the function expressive of the law of human mortality and on a new mode of determining life contingencies." *Philos. Trans. R. Soc. Lond.* 115:513–585.

Hall, Robert E. and Charles I. Jones. 2004. "The Value of Life and the Rise in Health Spending." *NBER Working Paper* 10737.

Hubbard, R. Glenn, Jonathan Skinner and Stephen P. Zeldes. 1994. "The importance of precautionary motives in explaining individual and aggregate saving." *Carnegie-Rochester Conference Series on Public Policy* 40:59–125.

Human Mortality Database. 2006. University of California, Berkeley (USA) and Max Planck Institute for Demographic Research (Germany). Available at www.mortality.org.

Hurd, Michael D. 1987. "Savings of the Elderly and Desired Bequests." *American Economic Review* 77(3):298–312.

Hurd, Michael D. 1989. "Mortality Risks and Bequests." *Econometrica* 57(4):779–813.

Johansson, Per-Olov. 2002. "On the Definition and Age-Dependency of the Value of a Statistical Life." *Journal of Risk and Uncertainty* 25(3):251–263.

Kalemli-Ozcan, Sebnem, Harl E. Ryder and David N. Weil. 2000. "Mortality decline, human capital investment, and economic growth." *Journal of Development Economics* 62:1–23.

Lee, Ronald D. and Lawrence R. Carter. 1992. "Modeling and Forecasting U.S. Mortality." *Journal of the American Statistical Association* 87(419):659–671.

Li, Qi and Shripad Tuljapurkar. 2004. "Life, Death, and the Economy: Mortality Change in Overlapping-Generations Models." *Michigan Retirement Research Center Working Paper* WP 2004-072.

Lucas, Robert E. 1987. *Models of Business Cycles*. New York: Blackwell.

Murphy, Kevin M. and Robert H. Topel. 2005. "The Value of Health and Longevity." *NBER Working Paper* 11405.

Nordhaus, William D. 2002. The Health of Nations: The Contribution of Improved Health to Living Standards. In *The Economic Value of Medical Research*, ed. Kevin M. Murphy and Robert H. Topel. Chicago: University of Chicago Press.

Rosen, Sherwin. 1988. "The Value of Changes in Life Expectancy." *Journal of Risk and Uncertainty* 1(3):285–304.

Ruhm, Christopher J. 2000. "Are Recessions Good for Your Health?" *Quarterly Journal of Economics* 115(2):617–650.

Ruhm, Christopher J. 2003. "Good Times Make You Sick." *Journal of Health Economics* 22(3):637–658.

Ruhm, Christopher J. 2004. "Macroeconomic Conditions, Health and Mortality." *NBER Working Paper* 11007.

Smith, James P. 1999. "Healthy Bodies and Thick Wallets: The Dual Relation between Health and Economic Status." *Journal of Economic Perspectives* 13(2):145–166.

Tapia Granados, José A. 2005. "Increasing Mortality During the Expansions of the US Economy, 1900–1996." *International Journal of Epidemiology* 10.1093/ije/dyi141.

Tolley, George S., Donald Scott Kenkel and Robert G. Fabian, eds. 1994. *Valuing Health for Policy: An Economic Approach.* Chicago: University of Chicago Press.

Tuljapurkar, Shripad and Ryan D. Edwards. 2005. "Variance in Death and Mortality Decline." Unpublished manuscript.

Vaupel, James W. and Vladimir Canudas Romo. 2003. "Decomposing Change in Life Expectancy: A Bouquet of Formulas in Honor of Nathan Keyfitz's 90th Birthday." *Demography* 40(2):201–216.

Viscusi, W. Kip. 1993. "The Value of Risks to Life and Health." *Journal of Economic Literature* 31(4):1912–1946.

Viscusi, W. Kip and Joseph E. Aldy. 2003. "The Value of a Statistical Life: A Critical Review of Market Estimates throughout the World." *Journal of Risk and Uncertainty* 27(1):5–76.

Figure 1: The distribution of life spans, survivorship, and log mortality in the U.S. in 1999

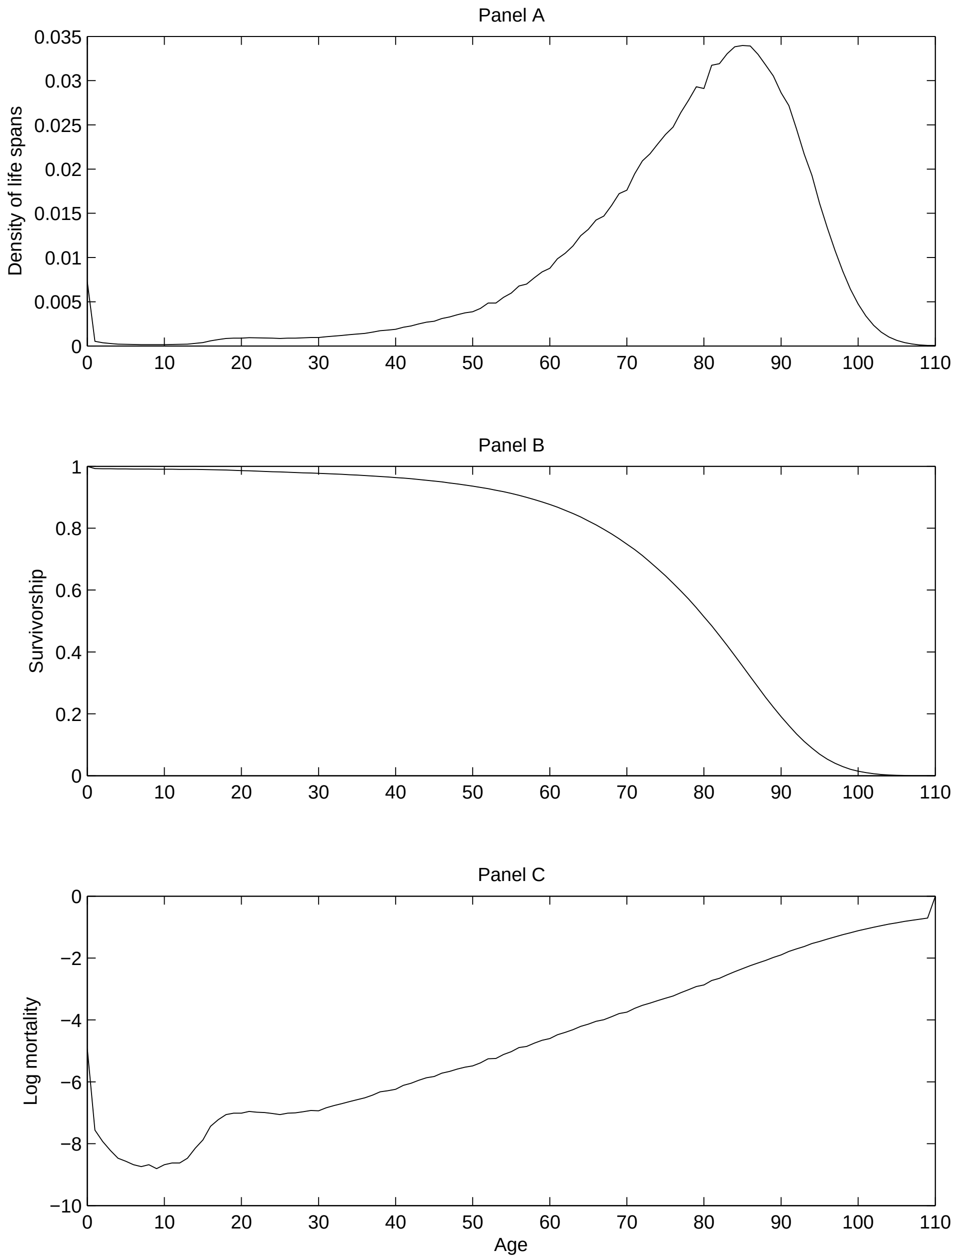

Notes: All data are provided by the Human Mortality Database (2006).

Figure 2: The price of S, a standard deviation in life span, in terms of mean life span as a function of δ when $r = 0.03$ and life span is normally distributed

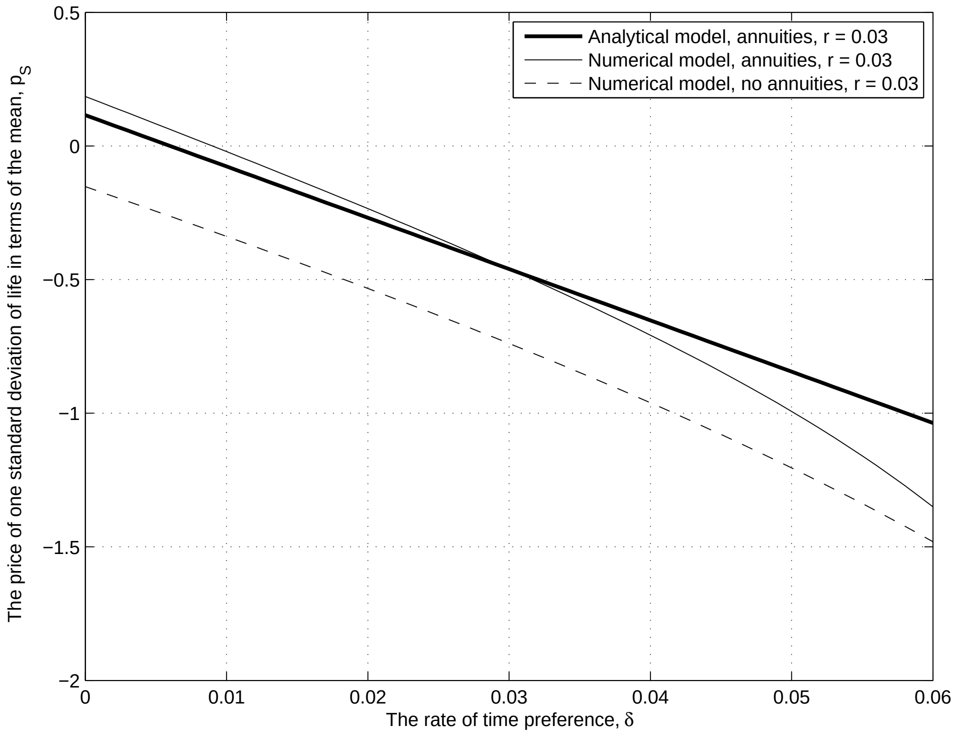

Notes: The thick black line shows p_S, the price of a standard deviation in life span in terms of mean life span, as a function of δ for the simple analytical model when the standard deviation of life span, $S = 15.66$, the level prevailing in the U.S. in 1994, and when $r = 0.03$. In the simple analytical model, when consumption is fixed, $p_S = -\delta S$. The thin solid line shows the locus in the full numerical model with complete annuitization of wealth. The dashed line depicts the locus in the model with no annuities.

Figure 3: The price of S, a standard deviation in life span, in terms of mean life span as a function of r when $\delta = 0.03$ and life span is normally distributed

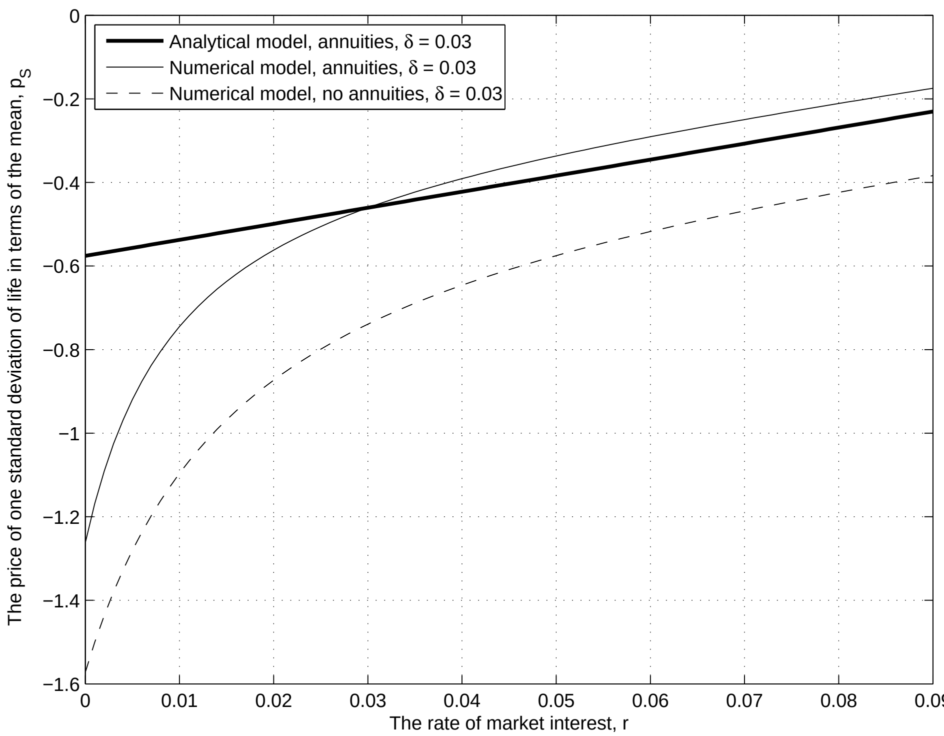

Notes: The thick black line shows p_S, the price of a standard deviation in life span in terms of mean life span, as a function of r for the simple analytical model when the standard deviation of life span, $S = 15.66$, the level prevailing in the U.S. in 1994, and when $\delta = 0.03$. In the simple analytical model, when consumption is fixed, $p_S = -\delta S$. The thin solid line shows the locus in the full numerical model with complete annuitization of wealth. The dashed line depicts the locus in the model with no annuities.

Figure 4: The price of S, a standard deviation in life span, in terms of mean life span as a function of δ when $r = 0.03$, adult survivorship is realistic, but there is no infant mortality

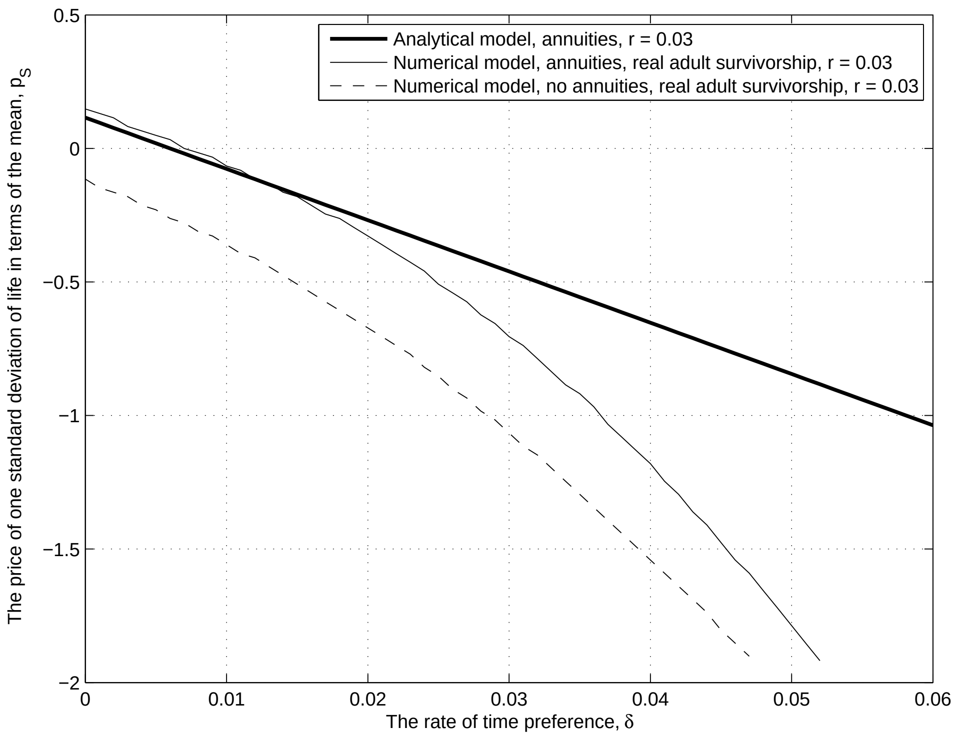

Notes: The thick black line shows p_S, the price of a standard deviation in life span in terms of mean life span, as a function of r for the simple analytical model when the standard deviation of life span, $S = 15.66$, the level prevailing in the U.S. in 1994, and when $\delta = 0.03$. In the simple analytical model, when consumption is fixed, $p_S = -\delta S$. The thin solid line shows the locus in the full numerical model with complete annuitization of wealth. The dashed line depicts the locus in the model with no annuities. In this simulation, the survivorship weights are taken from a modified period life table for the U.S. in 1999, provided by the Human Mortality Database (2006). Life-table deaths at ages under 10 are set to equal deaths at age 10, and the entire distribution is rescaled to sum to unity.

Figure 5: The price of S, a standard deviation in life span, in terms of mean life span as a function of δ when $r = 0.03$ and survivorship is fully realistic, with infant mortality

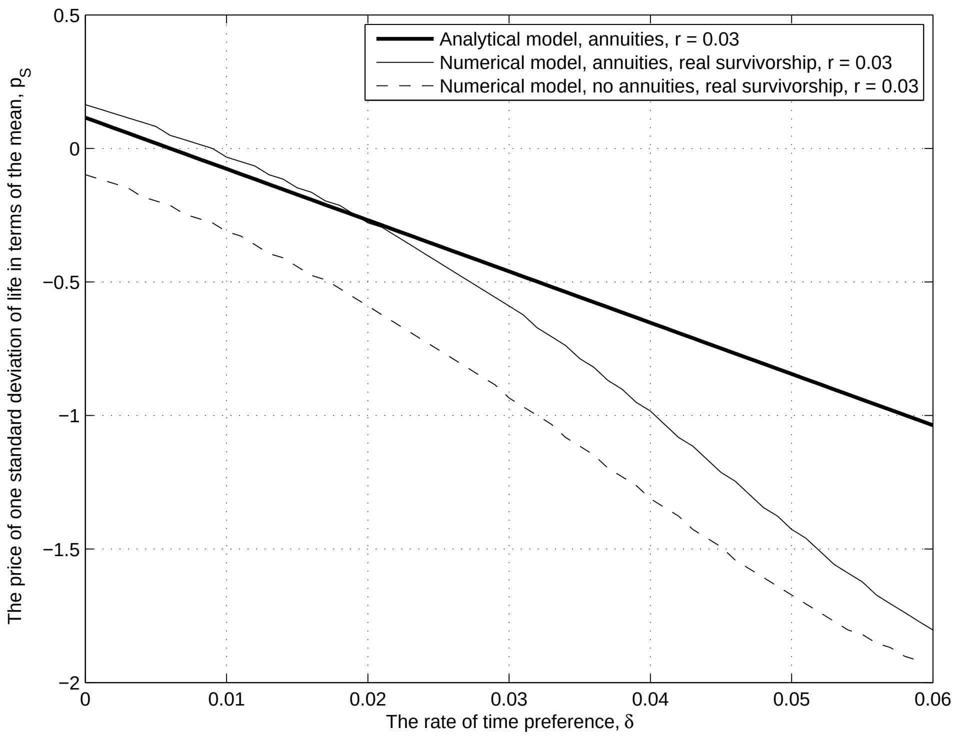

Notes: The thick black line shows p_S, the price of a standard deviation in life span in terms of mean life span, as a function of r for the simple analytical model when the standard deviation of life span, $S = 15.66$, the level prevailing in the U.S. in 1994, and when $\delta = 0.03$. In the simple analytical model, when consumption is fixed, $p_S = -\delta S$. The thin solid line shows the locus in the full numerical model with complete annuitization of wealth. The dashed line depicts the locus in the model with no annuities. In this simulation, the survivorship weights are taken from a modified period life table for the U.S. in 1999, provided by the Human Mortality Database (2006). Life-table deaths at ages under 10 are set to equal deaths at age 10, and the entire distribution is rescaled to sum to unity.

Figure 6: The standard deviation in life span above age 10 in the U.S. by year and by birth cohort

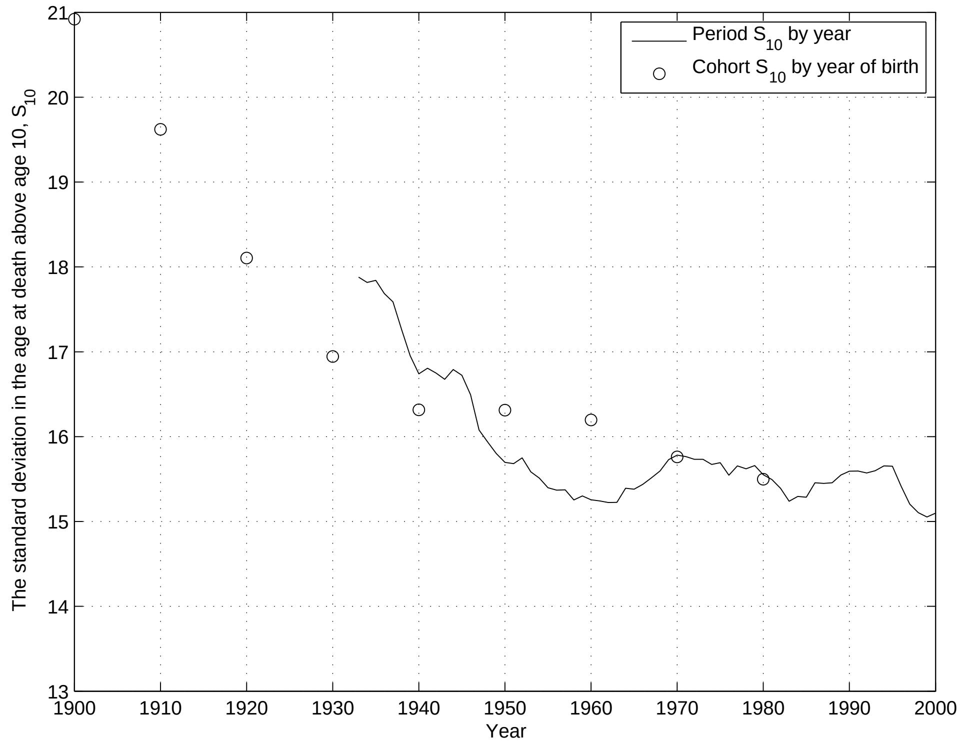

Notes: The underlying data are historical and forecast life table death distributions taken from Bell and Miller (2005). The standard deviation in life span above age 10, S_{10}, is calculated as described by Edwards and Tuljapurkar (2005).