The Cost of Uncertain Life Span*

Ryan D. Edwards†

January 18, 2012

Abstract

A considerable amount of uncertainty surrounds the length of human life. The standard deviation in adult life span is about 15 years in the U.S., and theory and evidence suggest it is costly. I calibrate a utility-theoretic model of preferences over length of life and show that one fewer year in standard deviation is worth about half a mean life year. Differences in the standard deviation exacerbate cross-sectional differences in life expectancy between the U.S. and other industrialized countries, between rich and poor countries, and among poor countries. Accounting for the cost of life-span variance also appears to amplify recently discovered patterns of convergence in world average human well-being. This is partly for methodological reasons and partly because unconditional variance in human length of life, primarily the component due to infant mortality, has exhibited even more convergence than life expectancy.

Keywords Health inequality Uncertainty Population health Convergence
JEL Classification I10 · J17 · O11

1 Introduction

Over the last two centuries, the demographic transition has produced vast reductions in human mortality rates and roughly a doubling in the average length of life among advanced countries (Lee, 2003). Since the 1950s, period life expectancy at birth, e_0, has exhibited much convergence among developed nations, and the average across those countries has increased at a roughly linear annual rate (White, 2002). The maximum female life expectancy at birth has exhibited similar trends over a much longer period (Oeppen and Vaupel, 2002). In the developing world, progress against mortality was delayed until the twentieth century, and in some regions it has been at least partially reversed by the

*This research was partially supported by NICHD grant T32 HD 07329.

†Associate Professor of Economics, Queens College and the Graduate Center, City University of New York, and NBER. I thank Deborah Balk, Neil Bennett, David Canning, Victor Fuchs, Ronald Lee, Shripad Tuljapurkar, David Weil, seminar participants at Queens College, CUNY, and at the Harvard Initiative for Global Health, and several anonymous referees for helpful comments.

HIV/AIDS epidemic or the fall of communism (Wilson, 2001; Deaton, 2004; Cutler, Deaton and Lleras-Muney, 2006). But there is also much evidence of convergence in demographic conditions, particularly when the focus is period life expectancy at birth (Wilson, 2001; Becker, Philipson and Soares, 2005).

The mean of the distribution of length of life is arguably the most informative, but the shape of the distribution is also interesting and, as I argue in this paper, economically important. The three panels in Figure 1 depict three representations of the same underlying period mortality data for the U.S. in 1900 and 2000, between which life expectancy at birth grew from 47.7 to 76.7 years. Probability distributions of ages at death appear in panel A, survivorship probabilities are in panel B, and age-specific log mortality rates are shown in panel C.[1] All three panels reveal a large amount of historical shape-shifting, but the dynamics in the distribution of length of life, shown in panel A, are the primary subject of this paper. The distribution moved rightward, as one would expect given the increase of 29 years in e_0 over this period. But a very large part of this improvement was due to the massive reduction in infant mortality that is visible in the graph. For individuals surviving infancy and childhood, past say the age of 10, the mean length of adult life, which I denote M_{10}, rose by a more modest 17 years, from 60.6 to 77.4, while the modal adult age at death increased from 72 to 85.[2] As the old-age mode moved rightward, the amount of variance around that mode fell notably. The standard deviation in length of life above age 10, or S_{10}, fell 5.3 years, from 20.0 to 14.7. These standard deviations in adult life span, a measure introduced by Edwards and Tuljapurkar (2005) that will become useful later, are superimposed in Figure 1A as horizontal bars extending on either side of the adult mean.[3]

Once we examine the shape of the distribution rather than just its mean, the simple story of broad cross-country convergence in mortality becomes more complicated. Edwards and Tuljapurkar (2005) reveal large differences among industrialized countries since 1960 in levels and trends in the variance of adult age at death as measured by S_{10}. While e_0 may have converged among these countries (White, 2002), the variance accounts for an increasing amount of mortality divergence between them. Separately, Edwards (2011) examines length of life distributions in a broad cross section of rich and poor countries in 1970 and

[1]The data are simple averages of sex-specific period life tables supplied by Bell and Miller (2005). Statistics prior to 1933, when the NCHS Death Registration Area expanded to include all states at 90 percent completeness, are of lower quality and include interpolated values.

[2]Remaining life expectancy, e_x, which is the entry commonly seen in a life table, is equal to average length of life conditional on survival to age x, M_x, minus x: $e_x = M_x - x$. It is useful to calculate M_x because of its role in the standard deviation and other central moments.

[3]Edwards and Tuljapurkar (2005) argue that the standard deviation in life span after age 10, S_{10}, is a good measure of this dispersion in adult life span. Age 10 is a convenient cutoff because it facilitates the examination of longer historical series or wider cross sections in which premature mortality is more prevalent. Opinions vary regarding the ideal type of variability measure for length of life. Wilmoth and Horiuchi (1999) and Fuchs and Ersner-Hershfield (2008) prefer the interquartile range, while Shkolnikov, Andreev and Begun (2003) favor the Gini, and Smits and Monden (2009) prefer the Theil index measured above age 15. Preferable qualities of S_{10} are that it is invariant to trends in infant mortality, which is etiologically distinct from adult mortality, and that like the mean, it is analytical tractable.

Figure 1: Distributions of life span, survivorship, and log mortality in the U.S. in 1900 and 2000

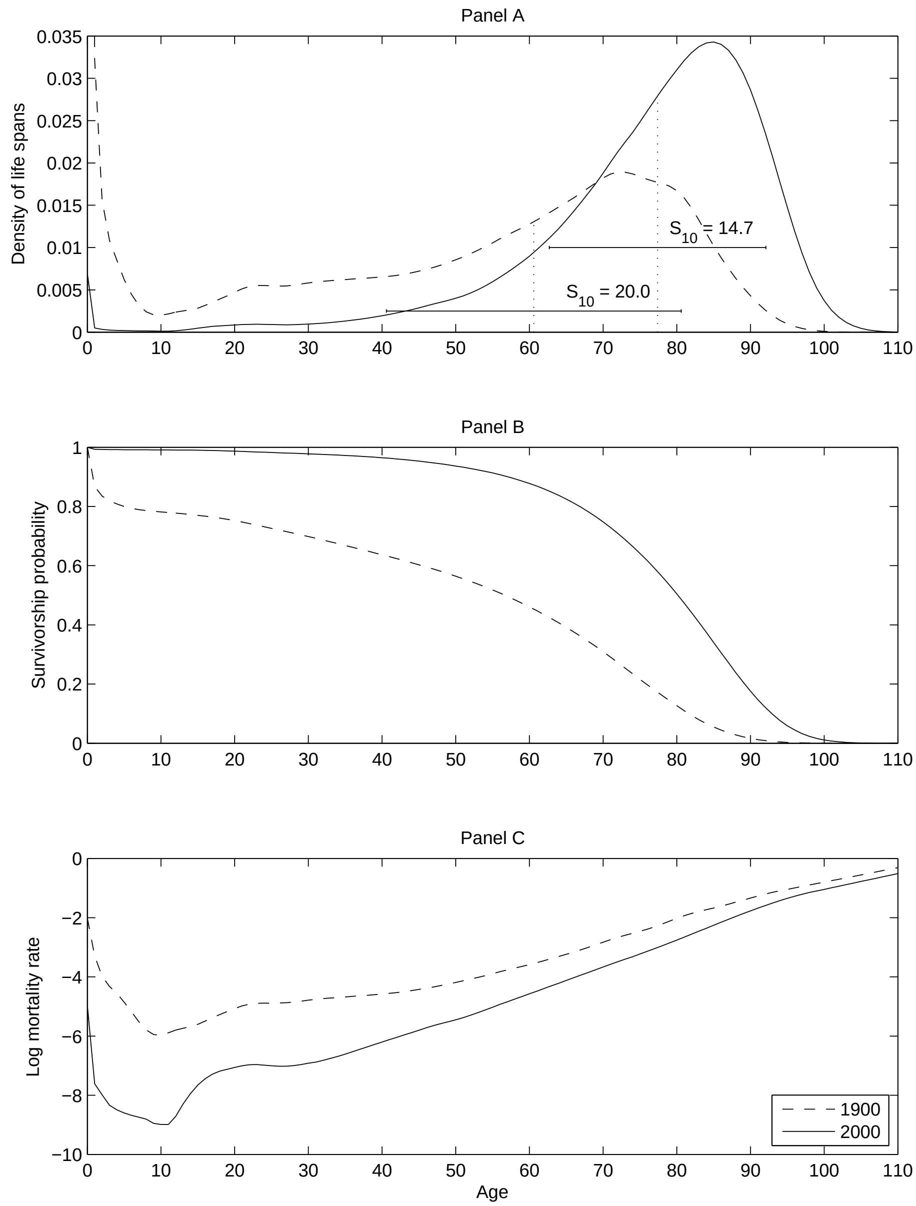

Data sources are period life tables from Bell and Miller (2005) and author's calculations, and they are unweighted averages of sex-specific life-table entries. In panel A, the density of deaths at age 0 was 0.1328 in 1900 and is left off the chart for expositional clarity. The horizontal bars in panel A depict one standard deviation above age 10, S_{10}, on either side of the mean above age 10, M_{10}.

2000, reporting evidence of declining within-country inequality, or the average level of S_{10}, but also indications of rising between-country inequality, especially in length of adult life. Infant mortality has exhibited much convergence across countries and has substantially narrowed differentials in e_0, the unconditional average length of human life (Wilson, 2001; Moser, Shkolnikov and Leon, 2005). But that narrative omits several interesting dynamics: relatively broad-based reductions in S_{10} among many developing countries owing to the demographic transition, a new divergence in S_{10} among industrialized countries, and widening differences in length of adult life across development groups.

What, if any, are the implications of these less obvious trends in mortality for human well-being? Are there welfare effects associated with S_{10} and infant mortality above and beyond the benefits reflected in e_0? In this paper, I show that the answer to the second question is yes, and that valuing longevity or life extension based on the mean alone, without accounting for the shape of the longevity distribution, typically leads to downward bias in estimates of the benefits of mortality decline. This is because the basic model of intertemporal choice over a finite lifetime, developed by Yaari (1965), implies that individuals are averse to risk or inequality over length of life under standard parameter assumptions. The basic intuition behind this result proceeds as follows. Consider a consumer whose subjective discount rate equals the real interest rate. If perfect annuities are available, the individual will plan constant consumption and thus constant flow utility in each period of life. Expected lifetime utility then equals a constant times the survivorship function shown in Figure 1B, a downward sloping concave curve. By Jensen's Inequality, a mean-preserving spread in length of life reduces expected lifetime utility in such a scenario.

Risk and inequality are conceptually related in choice theory (Atkinson, 1970), but they are not strictly the same in this context. The total variance in length of life in the population, as shown in Figure 1A, consists of between-group and within-group inequality. The former includes well-known socioeconomic gradients in life expectancy (Kitagawa and Hauser, 1973), while the latter would include true uncertainty faced by individuals who know their socioeconomic status and other characteristics that are important for longevity. But while differences between groups are certainly important, within-group inequality in length of life tends to be large (Edwards and Tuljapurkar, 2005). In this paper I will parsimoniously treat variance in length of life as though it were equivalent to uncertainty, but in reality different groups face different levels of uncertainty.[4]

For typical parameter values, the standard intertemporal model of Yaari (1965) suggests that variance in adult life span is costly, and that infant mortality is extremely costly, considerably more so than its effect on e_0 alone implies. In industrialized countries, each year in adult standard deviation is currently worth about half a year in the average; that is, an individual would agree to

[4] My approach is similar to those of Jones and Klenow (2010), who price consumption inequality using a Rawlsian perspective, or before any correlates with income are known, and to Lucas (1987), who prices aggregate consumption uncertainty without regard to heterogeneity in the variance of consumption across consumers. In a closely related paper (Edwards, 2009), I consider a problem that is analogous to Lucas's: the cost of cyclical fluctuations in mortality.

give up half a year in mean life span to obtain a standard deviation that is one year lower. This result lends new salience to the persistent differences we see across advanced countries today in their levels of S_{10}, which can be as large as 2 years. My results suggest that reductions in S_{10} in the U.S. were worth perhaps 9 percent of the total value of gains against mortality prior to 1950 but only 5 percent since then, because progress against S_{10} has slowed significantly. I also find that explicitly accounting for infant mortality and S_{10} amplifies the global convergence results of Becker, Philipson and Soares (2005), who priced the value of gains in income and e_0 between 1960 and 2000. Further, it might partially reverse some of the findings of Jones and Klenow (2010), who build on the work by Becker, Philipson and Soares by including consumption inequality and leisure in their measure of consumption-equivalent welfare. Jones and Klenow also account only for e_0 and report reduced welfare in Africa and Latin America between 1980 and 2000.[5]

More broadly, my results suggest that because variance in length of life is costly, medical decision making should probably take it into account in addition to average length of life. Increased interest in end-of-life care, which must balance patients' wishes, quality of life, and the potential reward of risky medical procedures, is a manifestation of a similar perspective (Kwok et al., 2011; Kelley, 2011). For populations, reducing aggregate health inequality takes on additional salience given the relatively high costs of variance in life span compared to the benefits of longevity increases.

The structure of the paper is as follows. In section 2, I contextualize my work by discussing previous research on choice under uncertain survivorship, the value of life, and bequests. In addition to Lucas (1987), Becker, Philipson and Soares (2005), and Jones and Klenow (2010), my work is also related to recent research on preferences over periods of life by Bommier (2006, 2008), and Bommier and Villeneuve (2011). In section 3, I examine a simple utility-theoretic model of the cost of uncertain length of adult life based on the framework of Yaari (1965), and I recover a convenient analytic solution for the cost of S_{10} that compares favorably to solutions from numerical simulations. I also discuss the less tractable cost of infant mortality, which proves to be very important in assessing the value of longevity gains during the demographic transition. In section 4, I explore the implications of these insights for decomposing the value of mortality declines and assessing convergence in average human well-being. The final section offers some concluding remarks.

[5]Previous studies of the total economic value of mortality decline in the U.S. since 1900 such as Nordhaus (2003) and Murphy and Topel (2006) typically account for all changes in the survivorship schedule, whether involving the mean, variance, or higher moments, and thus provide unbiased results. But studies of international well-being such as Becker, Philipson and Soares (2005) and Jones and Klenow (2010) often use only e_0 due to data limitations and thus cannot account for the shape of the life-span distribution.

2 Background

2.1 The value of a statistical life and mortality reduction

This paper examines the welfare cost of variance in length of life, which depends critically on how consumers value life years. A robust literature explores the valuation of life and the willingness to pay for mortality reduction using theoretical and empirical perspectives (Rosen, 1988; Viscusi, 1993; Tolley, Kenkel and Fabian, 1994; Johansson, 2002; Viscusi and Aldy, 2003; Aldy and Viscusi, 2008). These papers typically assume that agents maximize expected utility, and studies recover empirical estimates of the value of a statistical life (VSL) by observing wage differentials associated with mortality risks that vary across occupations. Resulting estimates of the VSL life vary widely, as reported by Viscusi and Aldy (2003) in their meta-analysis, with a range of $4 to $9 million in the U.S. They also appear to vary systematically across time and space, probably because income and consumption are changing, but the literature does not agree on precisely how the VSL varies. Viscusi and Aldy (2003) report an income elasticity of the VSL of about 0.5 based on international cross sections, while Costa and Kahn (2004) recover an elasticity of about 1.6 in longitudinal U.S. data.

In theory, the value of a statistical life should be a function of the value of remaining life years. Researchers have combined this principle with assumptions about the shape of intertemporal preferences, typically the framework of additive separability introduced by Yaari (1965) and discussed below, in order to recover measures of the value of a statistical life year (VSLY) from estimates of the VSL derived from hedonic wage regressions. In its most simple and common form, this approach posits that the age slope of the VSLY is either flat, if the discount rate is zero, or downward sloping if it is positive. But it can also be hump-shaped when calibrated with life-cycle patterns of consumption and earnings (Murphy and Topel, 2006; Aldy and Viscusi, 2008). Using such an approach and a posited income elasticity of the VSL, several studies measure the consumption-equivalent value of long-term improvements in longevity (Cutler and Richardson, 1997; Nordhaus, 2003; Murphy and Topel, 2006), typically finding that the value of mortality decline rivals the value of economic growth. Others have combined the value of longevity extension with growth in income or consumption in order to assess convergence in total human well-being (Becker, Philipson and Soares, 2005; Jones and Klenow, 2010). Most of these efforts utilize realistic survivorship probabilities like those shown in panel B of Figure 1, and thus are accounting for changes in the shape of the life-span distribution such as depicted in panel A. But studies of international well-being typically cannot due to data limitations (Becker, Philipson and Soares, 2005; Jones and Klenow, 2010). As a result, these studies measure life span using e_0 alone and none of the higher moments.

2.2 Preferences over length of life

To be sure, higher moments do not matter if there is risk neutrality over length of life, and the degree of that risk aversion is not trivially clear and could be zero or even negative. This topic is particularly salient for the medical profession, where decisions regarding life and death and the costs, benefits, and riskiness of procedures must be weighed by physicians and patients alike. Assessing net benefits requires assumptions about preferences over health states in different future periods, and the medical literature recognizes that the degree of risk aversion over remaining years of life will affect this calculation (Ried, 1998; Bleichrodt and Quiggin, 1999). Measuring the concavity of preferences over life span is neither a common nor straightforward activity, but the consensus view based on empirical research in the last two decades seems to be that individuals are risk averse. Bleichrodt and Johannesson (1997) report that four out of five empirical studies directly examining this question reject risk neutrality over life years in favor of risk aversion (McNeil, Weichselbaum and Pauker, 1978; Stiggelbout et al., 1994; Verhoef, Haan and van Daal, 1994; Maas and Wakker, 1994). These investigations typically ask respondents, sampled either from medical treatment programs or from the community, to assess the desirability of various probabilistic scenarios regarding survival in perfect health versus death. For many but not all respondents, certainty equivalents are concave in life years, implying risk aversion.[6]

Do standard models of preferences based on the work of Yaari (1965) imply risk aversion over length of life? In the next section, I show that the answer is yes, given certain reasonable conditions. But my approach is not the first nor the last word on the topic, which is an interesting new area of research. In several recent contributions, Bommier (2006, 2008) and Bommier and Villeneuve (2011) develop and use a generalized version of the additively separable preferences introduced by Yaari (1965). By specifying a recursive utility function, they introduce an extra preference parameter that governs the curvature of expected utility over length of life in addition to the time discount rate, which already provides curvature under reasonable assumptions, as I will show in the next section. In contrast, Bommier (2006) argues that Yaari's framework even with discounting implies risk neutrality over length of life, but his definition of risk neutrality is zero utility curvature along a special life-cycle consumption path that may or may not be representative.[7] Because preferences over length of life depend on consumption in a mechanical way, Bommier defines risk at-

[6] The shape of preferences tends to vary by subgroup characteristics and based on whether the gamble is short or long-term in nature. Pliskin, Shepard and Weinstein (1980) reveal apparent risk neutrality and even risk preference among 10 Harvard researchers. Verhoef, Haan and van Daal (1994) find their subjects are risk-seeking over small gambles in life span but risk averse over large gambles, consistent with prospect theory. Miyamoto and Eraker (1985) settle on risk neutrality over life years as an average over the wide-ranging preferences they observe.

[7] A consumption trajectory fitting Bommier's definition of the constant flow of satisfaction consumption profile, where $e^{-\delta t}u(c(t))$ is time invariant given a constant time discount rate δ, must be increasing through age because marginal utility is positive. This does not appear to be typical (Banks, Blundell and Tanner, 1998).

titudes around a "constant flow of satisfaction consumption profile," in which discounted flow utility, which is also the marginal utility of an additional year of life, is constant over time. Under such circumstances, standard measures of risk aversion in length of life, which depend on the second derivative of utility, must be zero because the marginal utility of additional years of life is assumed to be constant. Viewed this way, Bommier's risk neutrality result seems tautological and far from general.

But it is clear that this new literature fills an obvious gap and raises some important questions. The contributions of Bommier (2006, 2008) and Bommier and Villeneuve (2011) are similar in methodology and motivation to those of Kreps and Porteus (1978) and Epstein and Zin (1989, 1991), who use recursive models to separate preferences over financial risk from intertemporal substitution, and their success in fitting empirical patterns may be similar. Bommier and Villeneuve (2011) find that additional risk aversion over length of life than is implied by time discounting alone provides a better fit to the empirical age trajectory in the value of a statistical life reported by Aldy and Viscusi (2003), which has an inverted-U shape. But Murphy and Topel (2006), Hall and Jones (2007) and others reproduce similar trajectories by modeling health in the utility function, or by allowing for quality adjustment of life years by a means other than injecting more risk aversion. As Bommier and Villeneuve remark, the debate over the degree of risk aversion in length of life, like the ongoing debate over risk aversion in financial decisions, is far from settled but is worth engaging. If risk aversion is indeed greater than implied by time discounting alone, then my estimates will understate the true cost of life-span uncertainty. But Nordhaus (2003), Becker, Philipson and Soares (2005), Murphy and Topel (2006), Hall and Jones (2007); Jones and Klenow (2010), and others in the subfield on valuing longevity extension all assume additive separability as posited by Yaari (1965). My contribution follows in their footsteps and highlights the independent and interesting cost of the variance in the widely used baseline model; it makes no progress in answering similarly interesting questions about nonseparable preferences in theory and practice.

The inverted-U-shaped age trajectory of the VSL and the VSLY in recent empirical studies is both controversial and independently interesting for this study. As Aldy and Viscusi (2008) discuss, the result contrasts with the basic intuition that newborns should value their lives most highly because they have the most years left to live. The simple model I propose in the next section incorporates that dynamic as a reasonable baseline. But if the true VSLY follows an inverted-U shape, then the cost of variance in adult length of life may be overstated in the simple model. This is because a mean-preserving spread increases old life years at the expense of young life years, which in the simple model are more dear because of discounting. But the reason why the VSLY rises before declining with age is arguably important in a welfare study. If the accumulation and depreciation of a broad version of human capital including knowledge, skills, and health is the driving force, then measuring a heightened loss of welfare with death at middle age seems appropriate. If instead the VSLY follows the hump-shaped earnings age profile because of binding liquidity constraints, as Aldy and Viscusi (2008) seem to suspect, then it is less clear whether a welfare study ought to price a market failure that probably reduces well-being though ideally it should not. My approach values young human life the most, a reasonable but not uncontroversial baseline.

2.3 Annuities and bequests

Even if individuals are inherently risk averse over life span, they may be able to diversify life-span risk through contingent claims. Indeed, observed heterogeneity in stated preferences may reflect differential access to or use of markets for contingent claims. Annuities and life insurance are two examples of market instruments that diversify risks associated with uncertain life spans, while bequests are a nonmarket instrument that could also conceivably hedge against life-span uncertainty. In this paper, I show that while full annuitization removes all *consumption risk* associated with uncertain life span, and therefore improves welfare, annuities cannot remove the *utility risk*. Even under full annuitization, life-span uncertainty is costly. Whether life insurance helps offset S_{10} is more difficult to say, since it affects individual utility only through the bequest motive. That is, actuarially fair life insurance is like a precommitted bequest, and to assess its benefits we must understand bequests.

If individuals are altruistic, it is conceivable that bequests could hedge life-span risk relatively well if the bequest motive is strong. With utility deriving solely from consumption rather than from other aspects of living, then the disutility of early death could in theory be balanced by increased utility among survivors under altruistic bequests. Similarly, the additional utility deriving from late death could be offset by the impact of diminished bequests. But it is difficult to see why such fully altruistic individuals with utility only from consumption would care about any moment of life span, including the average. Probably in part for this reason, the value-of-life literature typically ignores bequest motives altogether (Chang, 1991; Johansson, 2002).

In any event, the literature on bequests is mixed with regard to the strength of the motive, with some research indicating they are generally not intended (Hurd, 1987, 1989) and other research suggesting otherwise (Kopczuk and Lupton, 2007). A prevailing view in economics is that bequests are simply unused precautionary savings (Dynan, Skinner and Zeldes, 2004). Findings in the medical literature of risk aversion over length of life certainly suggest that bequest motives are either not universal or not strong enough to hedge against the risk of death. Another perspective on bequests is that they can be strategic, a quid pro quo promised in exchange for elderly care (Bernheim, Shleifer and Summers, 1985). Leaving aside the problem that living too long risks depleting bequeathable wealth as well as requiring informal care, we might interpret strategic bequests as merely another form of annuitization, if they are set aside up front.

In the next section, I explore the theoretical cost of uncertain life span in the standard intertemporal model of Yaari (1965), with special attention paid to the role of annuitization. I find that individuals who discount their future well-being in the standard way should be risk averse over life span, even when they can fully annuitize. Consistent with the value-of-life literature, I do not model a bequest motive explicitly. I express the cost of variance in life span in terms of the value of the mean, which probably lessens the sensitivity of my results to the assumption of no intended bequests. A bequest motive would work to reduce both the value of the mean length of life and the cost of variance simultaneously.

3 Modeling the cost of life-span uncertainty

In this section, I show how the basic intertemporal model of Yaari (1965) with time discounting and additive separability implies there is a large welfare cost associated with adult life-span uncertainty for reasonable parameter values. There is also a very high cost attached to infant mortality.

3.1 Setup of the model

Consider an expected utility maximizer at time $t = 0$ with an implicit rate of time discounting equal to δ and no bequest motive.[8] Following Yaari (1965), lifetime expected utility is the discounted 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, $\ell(t)$:

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

The survivorship function, $\ell(t)$, shown in panel B of Figure 1, is one minus the cumulative density function of life span, which is shown in panel A.

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, education, capital, or financial risk in this model.[9] In a market without annuities, the

[8] Yaari (1965) initially posits a more general form of the subjective discount function, but time discounting must be exponential at a constant rate in order to insure preferences are time consistent (Strotz, 1956). Bommier (2006) discusses preferences over length of life in other frameworks, such as with hyperbolic discounting and other types of time inconsistent preferences.

[9] 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 cost 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 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. But it is also true that S_{10} is lower for groups with more education (Edwards and Tuljapurkar, 2005). We might interpret this as very tangible evidence that S_{10} is costly.

budget constraint requires the individual to finance the present value of all future consumption out of wealth:

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

Under this budget constraint, the model will produce unintended bequests whose size varies inversely with length of life. If instead actuarially fair annuities are available, the budget constraint takes a different form:

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

An annuity pays off in future periods only if the individual is alive. This allows the buyer to finance future consumption more cheaply than through saving, but at the expense of unintended bequests.[10]

The individual maximizes equation (1) subject either to (2) or (3) depending on whether annuities are available. The Euler condition that describes intertemporal choice is

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

where D is an indicator of the lack of annuities, and it multiplies $q(t)$, the mortality rate.[11] Under full annuitization, $D = 0$ and mortality cancels out of equation (4) because survivorship appears in both the objective and the constraint, and consumption will be flat through age if $\delta = r$. But without annuities, $D = 1$, and the mortality rate, $q(t)$, which typically increases exponentially through age, will pull marginal utility higher and consumption lower over age through a type of precautionary saving (Hubbard and Judd, 1987), producing a consumption trajectory that looks like the survivorship curve.

My analytical strategy is to assume full annuitization, normally distributed life spans, and power utility in order to find a convenient closed-form solution for the cost of life-span variance. Later, I employ numerical simulations to assess its fit. Because flow utility is a function of consumption, which typically depends on age and thus the life-span distribution, it is convenient to assume a specific functional form. Power utility is a standard assumption and a reasonable baseline.

[10] 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.

[11] In a discrete-time life table, $q(t)$ is the probability of dying between t and $t+1$, equal to $q(t) = \log[\ell(t)] - \log[\ell(t+1)]$ or 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.

3.2 Full annuitization and normally distributed life spans

Let $u(\cdot)$ take the familiar form of power utility with constant relative risk aversion over consumption plus a constant utility shifter K.[12]

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

Under full annuitization, survivorship weights appear both in lifetime utility and in the budget constraint. Thus mortality cancels out of the Euler equation (4), which simplifies to

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

where $c(0)$ is a function of wealth, the parameters, the annuitization indicator, D, and the moments of life span. I will proceed by assuming that $c(0)$ remains constant over small changes in the moments of life span, an assumption that I later relax in numerical simulations.[13]

With the simplified Euler equation, I can completely solve the model by reformulating the budget constraint through a change in the order of integration:

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

where T is a random variable, the realization of life span, and E is the expectations operator. The survivorship weights are now implicit in the expectation. For a given distribution of T, I could use equation (6) to solve the integral and

[12] To be sure, the use of power (isoelastic) utility raises some issues in any setting when periods of life are variable. Here, the level of utility matters, which is usually not the case. It must be positive or else life is not a good. We can model this with certain combinations of the utility shifter K and the coefficient of relative risk aversion, γ, as do Ehrlich (2000), Becker, Philipson and Soares (2005), and Hall and Jones (2007), but questions may remain as to the appropriateness of this technique. No study has attempted to test the restrictions imposed on the characteristics of $u(\cdot)$ in this case, namely the implied degree of risk aversion and intertemporal substitution in consumption, and analogous preferences over years of life. We already know power utility does not jointly satisfy the first two particularly well (Epstein and Zin, 1989, 1991), but that has not precluded its widespread use. To what extent modeling preferences over length of life complicate this picture is a question awaiting future research. I believe my core results hinge on the assumptions of time separability and exponential *time discounting*, and not on the curvature of the period utility function, or on the restriction that the coefficient of relative risk aversion in consumption equals the intertemporal elasticity of substitution in consumption.

[13] Mean life span, M, affects $c(0)$ in obvious ways, and Rosen (1988) shows how extending life incurs a marginal cost associated with reducing consumption in all periods, holding other things, namely liftime wealth, equal. A more subtle point is that the variance in life span, S^2, also affects $c(0)$ for the same reason that variance affects lifetime expected utility. But the direction of the effect is counterintuitive. Through Jensen's Inequality, (expected) lifetime discounted consumption is lower when variance in life span is higher. At any given initial wealth, $c(0)$ can then be higher than under less variance while still satisying the budget constraint. Numerical simulations reveal that the effects on $c(0)$ of changing M or S^2 are small.

then find $c(0)$ by taking the expectation.[14] But I am primarily interested in the relative price of variance in life-span in this model, which is governed by its relative marginal utility. To proceed, I change the order of integration in lifetime expected utility, equation (1):

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

where as before, T is a random variable. With the power utility formulation in equation (5) and the consumption function in equation (6), expected lifetime utility under full annuitization is

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

where

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

When r is close to δ, I have $\hat{\delta} \approx \delta$; and $\hat{\delta} = \delta$ when either $r = \delta$ or $\gamma \to 1$.

3.2.1 Risk aversion over life span

Examination of equation (9) reveals that individuals are risk averse over life span in this model if the rate of time discounting, δ, is positive and not too different from the real interest rate, r.[15] 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 $r = \delta$. That is, absolute risk aversion in life span is roughly constant under reasonable assumptions.

One would expect individuals who are risk averse over life span to be hurt by uncertainty in life span, and this is more easily seen if I assume that length of life is normally distributed:[16] $T \sim N(M, S^2)$. In that case, expected lifetime

[14] Without annuities, the presence of $q(t)$ in the consumption function precludes analytical solutions because mortality increases exponentially with age.

[15] As I discussed earlier, Bommier (2006) instead argues that the framework of Yaari (1965) implies risk neutrality over length of life. But the risk neutrality result derives from Bommier's assumption of a "constant flow of satisfaction consumption profile," in which $e^{-\delta t}u(c(t))$ is time invariant. That condition requires flow utility and thus consumption to be growing over time. As shown by the consumption trajectory in equation (6), consumption growth requires that r be higher than δ. In particular, if $\gamma = 0.8$, as I assume later, then $c(t)$ must be growing five times as fast as δ to produce growth in $u(\cdot)$ equal to δ. This would require $r = 5\delta$, a very large difference. If γ is bigger, flow utility is more sharply concave, requiring an even larger increase in consumption to raise utility by a particular amount. By equation (6), the difference between r and δ would then have to be even larger to achieve a particular level of consumption growth.

[16] As shown by Figure 1, adult life spans are technically not normal, with both leftward skewness and leptokurtosis, or peakedness with fat tails, an indicator of different subgroup variances. Below, I show that numerical simulations show that normality reduces the cost of S as long as the discount rate is positive, so this assumption produces an underestimate of the true cost.

utility in this model is

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

by virtue of the properties of lognormality. Expected lifetime utility is a decreasing function of S, provided that $\delta > 0$, r and δ are not too dissimilar, and period utility is positive.[17]

3.2.2 Pricing the variance in life span

It is convenient to recover 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:

$$p_S = \frac{\partial EU/\partial S}{\partial EU/\partial M}. \tag{12}$$

When the utility shifter K is nonzero, this ratio is mathematically complicated.[18] A first-order Taylor expansion around $K = 0$ reveals

$$p_S \approx -\hat{\delta}S + \left(\frac{c(0)^{1-\gamma}}{(1-\gamma)}\right)^{-1}\left(\hat{\delta}-\delta\right)e^{(\hat{\delta}-\delta)M+(\delta^2-\hat{\delta}^2)S^2/2}SK. \tag{13}$$

When $K = 0$ or when $\hat{\delta} = \delta$, this reduces exactly to a parsimonious relationship:

$$p_S = -\hat{\delta}S. \tag{14}$$

The price of a standard deviation in life span, p_S, is negative when $\hat{\delta} > 0$ because S is a bad. An individual who faces higher variance in life span must be compensated by a longer mean life span. In addition, p_S increases linearly with the level of S, with the magnitude of the slope equal to $\hat{\delta}$, approximately the rate of time discounting. That is, the costliness of a standard deviation in life span in terms of the mean rises with the level of uncertainty. Mathematically speaking, this follows directly from the lognormality of lifetime utility. Intuitively, the

[17] 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 reduces the cost of S through the second piece of (11), but numerical simulations confirm this effect to be small and uninteresting.

[18] The constant utility shifter does not appreciably augment the insights to be gained. Were K to describe period utility alone, it would imply the same dynamics as when $r = \delta$, with only δ mattering for cost. When K is combined with flow utility from consumption, both numerator and denominator in (12) are weighted averages of the two pieces in equation (11). When the piece including K has more weight, the coefficient on S in p_S shifts closer to δ than $\hat{\delta}$.

willingness to bear additional risk falls with the level of risk because its marginal disutility rises.[19]

It is also convenient to write out the marginal utilities in the special case of $K = 0$:

$$p_S = \frac{\partial EU/\partial S}{\partial EU/\partial M} = \frac{-\hat{\delta}^2 S}{\hat{\delta}} \frac{\frac{c(0)^{1-\gamma}}{(1-\gamma)\hat{\delta}} e^{-\hat{\delta}M+\hat{\delta}^2 S^2/2}}{\frac{c(0)^{1-\gamma}}{(1-\gamma)\hat{\delta}} e^{-\hat{\delta}M+\hat{\delta}^2 S^2/2}} = \frac{-\hat{\delta}^2 S}{\hat{\delta}} = -\hat{\delta} S. \quad (15)$$

The end result is the same as in the Taylor expansion, but this formulation reveals a subtle point that will become important in section 4. The marginal utility of mean life years, the denominator, is decreasing in the mean but increasing in the variance, so that improvements in life expectancy of a given size are more valuable when variance is high. This somewhat counterintuitive result again derives from discounting. Gaining an additional year at the mean under full certainty is not worth as much as gaining an additional year in expectation when there are higher survivorship weights on earlier, less heavily discounted years, and lower weights on later years.

Is $p_S = -\hat{\delta} S$ high or low? It clearly depends on the level of the discount rate. If we choose $r = \delta = 0.03$, their standard values in calibration exercises (Hubbard, Skinner and Zeldes, 1994; Becker, Philipson and Soares, 2005), then at the current U.S. level of $S = 15$, which I measure using the adult variance S_{10}, we find that $p_S = -0.45$ year. That is, the average citizen would be willing to give up almost half a year in mean life span in order to obtain a standard deviation in life span that was one year lower. For now, I simply remark that this cost seems large, and later I provide some context for assessing the cost relative to levels of population health across time and space.

3.3 Numerical solutions of the full model

In order to examine the sensitivity of the analytical result in equation (14) to alternative assumptions, I set parameters to match those used by Becker, Philipson and Soares (2005): $r = \delta = 0.03$, $\gamma = 0.8$, and $K = -16.2$. Initial wealth equaling \$800,000 is consistent with the parameter values, U.S. life spans, and per capita consumption of \$26,650 per year. I also fully endogenize consumption. I begin by modeling life span as normally distributed and omitting infant mortality, and then I reintroduce realistic adult and infant mortality.[20] For con-

[19] We see the same type of behavior in financial markets, where returns on financial assets are also approximately lognormal, and risk premia tend to rise strongly with the riskiness of returns. According to data presented by Ibbotson Associates (2002), the standard deviation of the excess return on equities was 14 percent between 1948 and 1999, which demanded a risk premium of about 9 percent. Excess returns on corporate bonds had a standard deviation a little over half as large, 8.5 percent, but the risk premium on corporate bonds was much lower, only 1.3 percent, or about one seventh of the equity risk premium.

[20] I truncate these synthetic distributions at ages 0 and 150 and rescale so that their cdf's sum to unity. Age 150 is an unrealistic but convenient choice when life spans are normally distributed. The Human Mortality Database (2009) topcodes age at 110, and there are few documented individuals who have survived to that age. When life spans are normally dis-

venience, I set adult mean and variance equal to U.S. levels in 1994 and search for the decrease in M that compensates for the observed decrease in S.[21]

3.3.1 Normally distributed life spans

The two panels in Figure 2 plot p_S as given by equation (14), shown by the thick solid line, on the same axes with two other loci that I obtain from numerical simulation of the model with normally distributed life spans. The thin solid lines depict the numerical model with annuities, while the dashed line shows results without annuities. In the top panel, I fix $r = 0.03$ and examine how varying δ changes p_S, while in the bottom panel I fix $\delta = 0.03$ and vary r.

In both panels, the two solid lines are fairly similar, revealing limited differences between the analytical and numerical versions of the model with annuities, except at the extremes. Both panels reveal that when δ is small relative to $r = 0.03$, the price of variance p_S draws close to zero. This is especially clear at the upper left in panel A, where δ is less than 0.01 while $r = 0.03$. This is essentially the risk neutrality result found by Bommier (2006) when consumption growth is rapid enough to keep the marginal utility of length of life constant, all because r exceeds δ by a lot. When time discounting is very low relative to the rate of interest and intertemporal substitutability is high, it is optimal to consume more in the future. A mean-preserving spread in life span, which trades away earlier years for later years, could actually improve expected well-being for somebody with heavily back-loaded consumption.

The dashed lines in Figure 2 show that the cost of uncertain life span is higher without annuities. It is striking that annuitization removes only a little over one third of the total cost of life-span uncertainty under baseline parameter values. That is, in Yaari's (1965) framework, the direct utility cost of S is relatively more important than the cost of consumption uncertainty. The period utility curvature parameter, γ, can play a role here but is relatively uninteresting.[22]

---

tributed 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 under realistic survivorship because real life spans are skew-left.

[21] In 1994, $\mathrm{M}_{10} = 76.85$ and $\mathrm{S}_{10} = 15.66$, while in 1999, $M_{10} = 77.67$ and $S_{10} = 15.05$. It is convenient to use data from these two years because there is a relatively large difference in S_{10} but a small difference in M_{10}, which reduces the complexity of later simulations with fully realistic mortality.

[22] As I discussed in note 12, the tradition in the literature on valuing life extension is not to separate the coefficient of relative risk aversion (CRRA) from the elasticity of intertemporal substitution (EIS), although researchers in asset pricing often model preferences that way (Epstein and Zin, 1989, 1991). When longevity is endogenous and must be purchased, as in the growth model of Hall and Jones (2007), this could be problematic because the marginal rate of substitution between consumption and length of life ought to depend on the EIS, not the CRRA. In the present context, the conflation of the EIS and CRRA is not particularly important. The influence of γ is best interpreted here as deriving from its role as the inverse of the EIS. When $0 < \gamma < 1$, the consumer likes to substitute consumption between periods because marginal utility in any period remains high, and any difference between r and δ will be amplified and will generally affect p_S. But when $\gamma > 1$, there are weaker gains from intertemporal redistribution, and high or low interest rates do not greatly affect $\hat{\delta} \approx \delta$.

Figure 2: The price of S, a standard deviation in life span, in terms of mean life span when life span is normally distributed

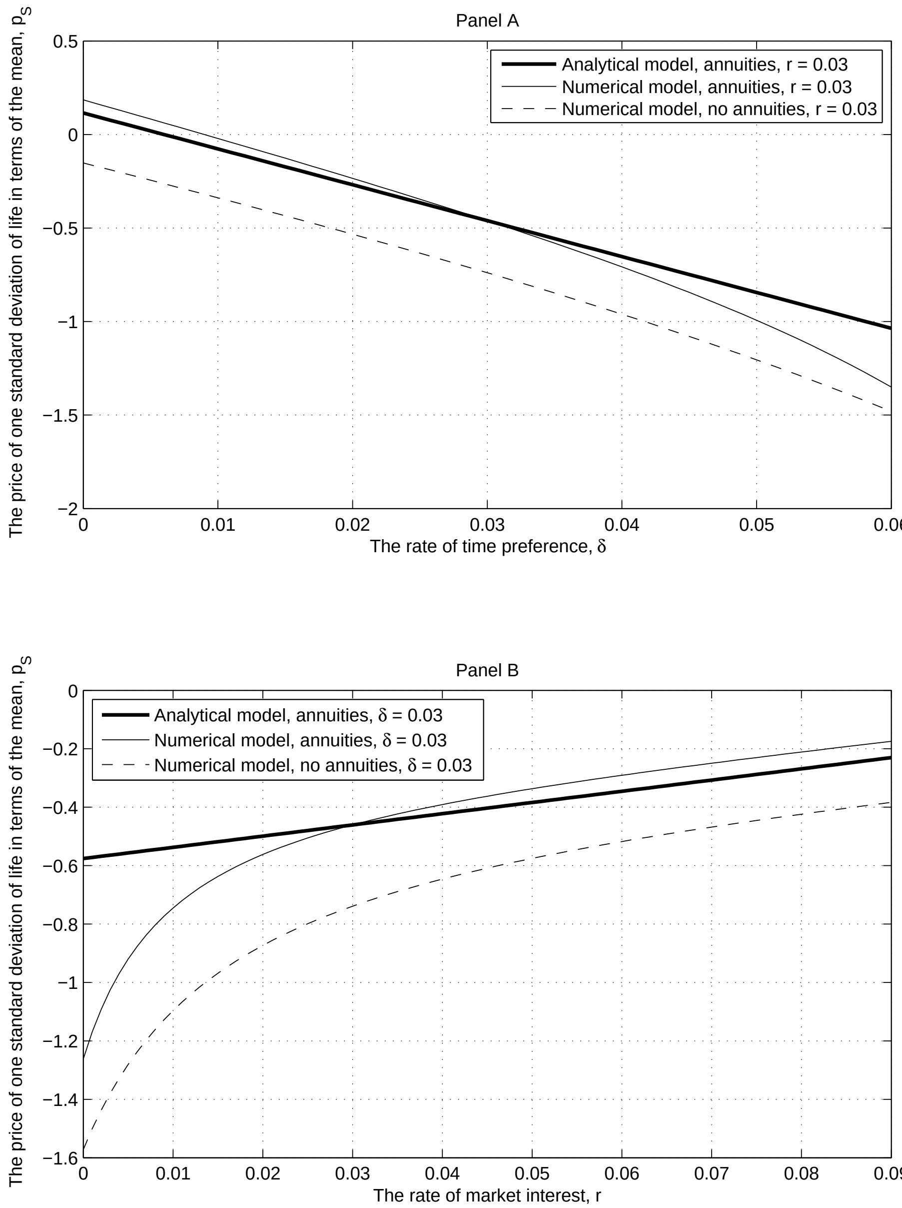

In both panels, the thick black lines show p_S, the price of a standard deviation in life span in terms of mean life span, either as a function of the rate of time preference, δ (panel A), or of the interest rate, r (panel B). The standard deviation of life span is set to $S_{10} = 15.66$, the level prevailing in the U.S. in 1994. In panel A, r is fixed at 0.03, and in panel B, δ is fixed at 0.03. The thin solid lines shows the locus in the numerical model with normally distributed length of life and complete annuitization of wealth. The dashed lines depicts the locus in a model with no annuities.

3.3.2 Realistic adult life spans without infant mortality

As shown in Figure 1A, modern distributions of human life span are skew-left and leptokurtic, not normal. Skewness implies that a mean-preserving spread in life span lowers survivorship probabilities asymmetrically; leftward skewness distributes an increase in variance more broadly below the mode or mean. Because of discounting, skewness should amplify the cost of uncertain life span. Figure 3A is the analogue of Figure 2A with realistic adult mortality.[23] It depicts the same three loci of p_S against δ for $r = 0.03$ and uses the same vertical scale for easier comparison. The thick solid line, which shows the analytical model's p_S, is the same in each, but the other two schedules are lower here, especially for larger δ. That is, the cost of life-span uncertainty is indeed higher under realistic adult mortality with leftward skewness.

3.3.3 Fully realistic life spans with infant mortality

Does the presence of infant mortality affect p_S? Given a fixed pattern of adult survivorship, higher infant mortality reduces the average length of life, raising the marginal utility of M. But if the marginal disutility of adult variance rises with total variance, infant mortality could amplify the marginal disutility of S. The effect on p_S is ambiguous. To proceed, I numerically simulate the cost of adult variance under fully realistic survivorship with infant mortality. I treat infant mortality as a completely separate dynamic, fixing infant deaths at their relative probability in 1999 and reestimate the compensating change in mean adult life span.[24] Panel B in Figure 3 depicts p_S as a function of δ with fully realistic mortality. Including infant mortality does not change the trajectories much, but it does slightly attenuate p_S compared to panel A. The presence of infant mortality must increase the marginal utility of mean life more than it

Functionally, the MRS between average length of life and the variance depends mostly on the time discount rate. Intuitively, it is utility curvature over length of life, not consumption, that matters most in the cost of life-span variance. The relevant issue here is whether additively time-separable preferences are overly confining, an issue discussed by Bommier (2006, 2008) and Bommier and Villeneuve (2011).

[23] It is tricky to model realistic life spans with particular means and variances because we do not have a convenient functional form of the actual probability distribution of life spans. 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 $\mathrm{S}_{10} = 15.66$. 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 preserving the variance. At ages under 10, I simply duplicate the density at 10 and renormalize the entire distribution. Later I include realistic infant mortality as explained in the text.

[24] 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. Then I renormalize so that each distribution cumulates to unity.

Figure 3: The price of S, a standard deviation in life span, in terms of mean life span with realistic mortality

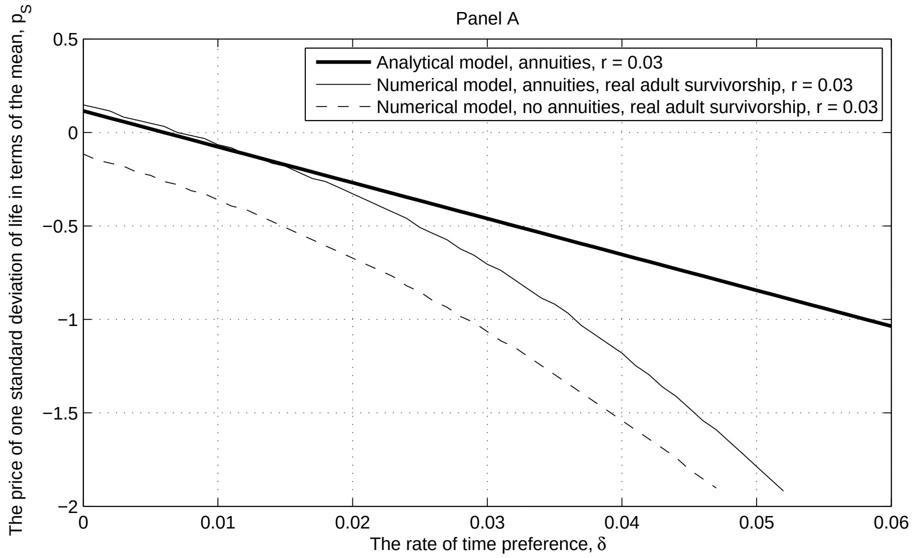

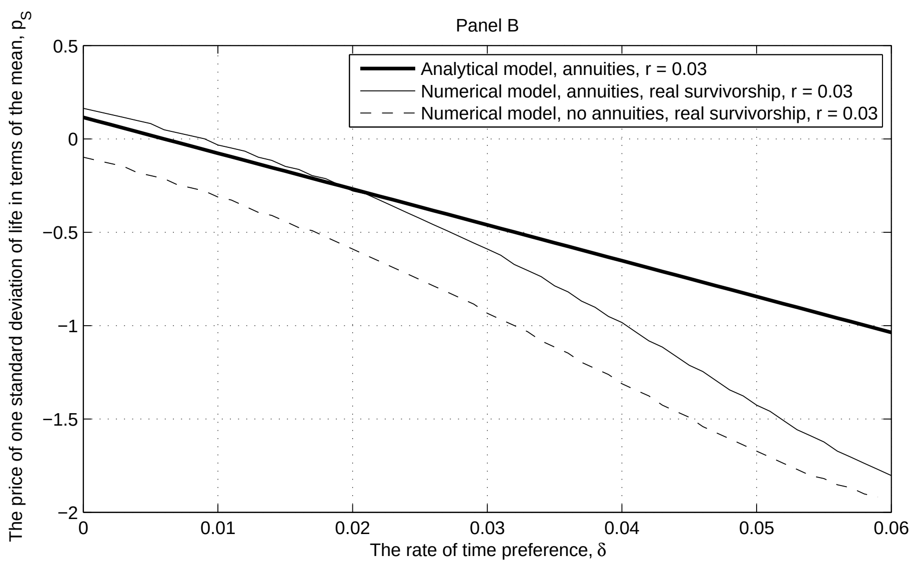

In both panels, the thick black lines show p_S, the price of a standard deviation in life span in terms of mean life span, as a function of the rate of time preference, δ, when r is fixed at 0.03. The standard deviation of life span is set to $S_{10} = 15.66$, the level prevailing in the U.S. in 1994. The thin solid lines shows the loci in the numerical model with annuities, while the dashed lines show them without. In panel A, the numerical model features realistic adult mortality and zero infant mortality, while in panel B mortality is fully realistic. Survivorship weights are from life tables provided by the Human Mortality Database (2009) and modified in order to reveal the results using a translation method described in the text. In panel A where infant mortality is omitted, 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. In panel B, the entire distribution is used, holding infant mortality constant as described in the text.

raises the marginal disutility of a standard deviation in life span.[25] But the bottom line is that the analytical solution remains a conservative estimate of the true cost.

3.4 The cost of infant mortality

A separate but related issue that arises with cross-country comparisons is the cost of infant mortality. The previous section shows it does not greatly affect p_S, and differences in infant mortality between advanced countries today are minor. But infant mortality is much higher and more volatile across time and space, and like S_{10}, it is costly because of the curvature of expected lifetime utility. To see why, assume for simplicity that infant mortality is completely described by $q(0)$, the probability of death between ages zero and one. A change in $q(0)$ equal to $i = -0.01$ raises survivorship, life expectancy at birth e_0, and lifetime expected utility each by 1 percent.[26] By comparison, proportional increases in e_0 brought about solely by reductions in adult mortality typically produce less than proportional increases in lifetime utility.[27] This is because discounting reduces the value of gains that are centered around the average length of life relative to the value of gains derived from infant mortality, which are broadly distributed across all ages. The implication is that declines in infant mortality are considerably more valuable than is implied by their effect on life expectancy at birth alone.

Formally, I can derive a formula for p_i, the cost of infant mortality in terms of the mean. Assuming complete annuities, normally distributed life spans, and $K = 0$ for simplicity,

$$p_i = \frac{\partial EU/\partial i}{\partial EU/\partial M} = -\frac{(1 - e^{-\hat{\delta}M + \hat{\delta}^2 S^2/2})/\hat{\delta}}{e^{-\hat{\delta}M + \hat{\delta}^2 S^2/2}} = -\frac{e^{\hat{\delta}M - \hat{\delta}^2 S^2/2} - 1}{\hat{\delta}}, \tag{16}$$

where flow utility cancels. Equation (16) states that a unit reduction in i is

[25] To assess how very high rates of infant mortality might change the marginal utility of mean life span and this p_S, I ran the same experiment using Swedish data from 1900, when deaths at age 0 were 10 percent. Results were very similar, with the analytical result still a conservative estimate so long as $\delta > 0.02$ when $r = 0.03$. If $\delta < 0.01$, p_S became positive with such high infant mortality.

[26] When schedule $q^*(t)$ differs from $q(t)$ only by $q^*(0) = q(0) + i$, survivorship is given by

$$\begin{aligned} \ell^*(t) &= \exp\left(-\int_0^\infty q^*(t)\ dt\right) = \exp\left(-[q(0) + i] - \int_1^\infty q(t)\ dt\right) \\ &= e^{-i} \exp\left(-\int_0^\infty q(t)\ dt\right) = e^{-i}\ell(t). \end{aligned}$$

Life expectancy at birth is the integral of survivorship, $e_0 = \int_0^\infty \ell(t)\ dt$, while lifetime expected utility is given by equation (1). With complete annuities and $r = \delta$, lifetime utility rises proportionally with i along with e_0 and $\ell(t)$.

[27] For example, if life expectancy $T = M$ with complete certainty, $r = \delta$, and consumption is fully annuitized, $EU = u(\bar{c})(1 - e^{-\delta M})/\delta$. When $\delta = 0.03$, an increase in M of one percent will raise EU by half a percent when M is around 40, by one-third of a percent when M is near 60, and by one-quarter of a percent when M is 80. Numerical simulation reveals that this relationship also holds when there is realistic but static variance in length of life.

preferable to a unit increase in M, which occurs when $p_i < -1$, if the compounded risk-adjusted return to investing at the discount rate for M years, the numerator, is higher than the discount rate, the one-period return. An interesting quantity is $p_i/100$, the price of one percentage point in the infant survival probability in terms of the mean length of life. By equation (16), when $\delta = 0.03$, $p_i/100$ equals one third of the risk-adjusted total return at δ over M years. The latter reaches and begins to exceed 1 in absolute value around $M = 40$ when there is realistic adult variance. Thus when life expectancy exceeds 40, as it typically does in human populations, a decline in infant mortality of a percentage point is more valuable than an additional year in the adult mean. A quick glance at Tables 1 and 2, which I discuss next, reveals that gains against infant mortality, loosely approximated by $1 - \ell_{10}$ or the probability of death before age 10, have been on the order of tens of percentage points last century.

4 Discussion and Extensions

4.1 The role of the discount rate

The key element in the model is the discount rate, δ, which is approximately the coefficient of absolute risk aversion over gambles in life span under reasonable assumptions. The discount rate is a latent preference parameter, but it is standard in the literature to set it equal to 3 percent, roughly the real rate of return on government bonds, and here I have followed suit. Viscusi and Aldy (2003) review estimates of the discount rate in the U.S. and report a very wide range of 1–17 percent, making it difficult to reject the hypothesis that market rates of interest and the discount rate are the same (Picone, Sloan and Taylor Jr., 2004). Via an evolutionary argument, Rogers (1994) suggests that the discount rate in human populations should equal roughly 2 percent in the long run, which is in the same ballpark.

Discount rates appear to vary over individuals within and across countries (Barsky et al., 1997; Becker and Mulligan, 1997; Bishai, 2004), as does length of life (Edwards and Tuljapurkar, 2005; Edwards, 2011). Complications may arise if discount rates and life spans are related to one another, but the implications are difficult to assess because the literature does not speak with one voice on the subject. Fuchs (1982) views the discount rate as determinining health investments, while Becker and Mulligan (1997) see wealth, uncertainty, and health or the length of life as jointly determining the discount rate. Bishai (2004) presents empirical evidence of reductions in the discount rate, i.e., increases in patience, with age and schooling. In related work, Satchell and Thorp (2011) argue that mortality and fertility patterns imply that family trusts ought to first increase and then decrease patience as it ages. At the population level, these concerns seem unlikely to be problematic for assessing the average cost of uncertain life span. But microeconomic implications of life-span variance may be very different if lines of causality running between the discount rate and mortality are important.

4.2 Uncertainty in actual life spans

Different groups and individuals face substantially different amount of life-span uncertainty, not all of which can be due to behavioral differences. Edwards and Tuljapurkar (2005) show that S_{10} is systematically lower by about 1 year among females compared with males, and that it is 2–3 years higher among African Americans relative to whites. Individuals in the lowest quintile of household income had 2.4 more years in standard deviation than those in the upper 80 percent, while those without a high school degree had 2.1 years more than high school graduates. A subgroup difference of 3 years in standard deviation implies an increase in the costliness of life-span uncertainty of 20 percent.[28]

With this much subgroup variation, one should be a little cautious about interpreting trends in aggregate uncertainty. An increase in aggregate S_{10} could reflect increasing between-group inequality, increasing within-group inequality, or both, and there are different implications of each. If an increase in life expectancy were enjoyed only by a more advantaged subgroup, it would technically increase aggregate S_{10} by raising between-group inequality, but it would also be Pareto-improving if the less advantaged did not lose. Although aggregate S_{10} in the U.S. has displayed little trend (Edwards and Tuljapurkar, 2005), recent evidence on widening educational differentials in life expectancy suggests between-group inequality is indeed increasing (Meara, Richards and Cutler, 2008). By implication, within-group inequality should have been decreasing, and if that were true, disadvantaged groups could be gaining from reductions in their S_{10} even while the rich are getting richer. These patterns pose interesting questions for future research.

What about members of different birth cohorts? Up to now, I have proxied the actual, or cohort life spans of individuals with those based on period mortality rates. Although period S_{10} is the appropriate variance analogue of period life expectancy, the most commonly cited population health statistic, we would also like to know the levels of variance faced by actual cohorts. Differences between cohort and period S_{10} do not appear to be large, as shown by Figure 4. I plot both measures for the U.S. since 1900 using decennial period life tables from Bell and Miller (2005) prior to 1930 and annual life tables from the Human Mortality Database (2009) starting in 1933, and decennial cohort life tables from Bell and Miller. The two series track each other relatively well, with both showing the enormous impact of the epidemiological transition early last century. Cohorts alive today, who likely face $\mathrm{S}_{10} = 15$, face drastically less uncertainty in their life spans than those born around 1900, for whom $\mathrm{S}_{10} = 21$.[29]

[28] The less fortunate would also bear a heavier burden if they have disproportionately less access to annuities, which I found 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 are also poor for liquidity constrained individuals.

[29] Technically, true cohort S_{10} should also reflect uncertainty about future mortality rates, or in other words, uncertainty about the shape of the probability distribution itself. We can treat forecast uncertainty as independent from what we might call life-table uncertainty,

Figure 4: The standard deviation in life span above age 10 in the U.S. by year and by birth cohort

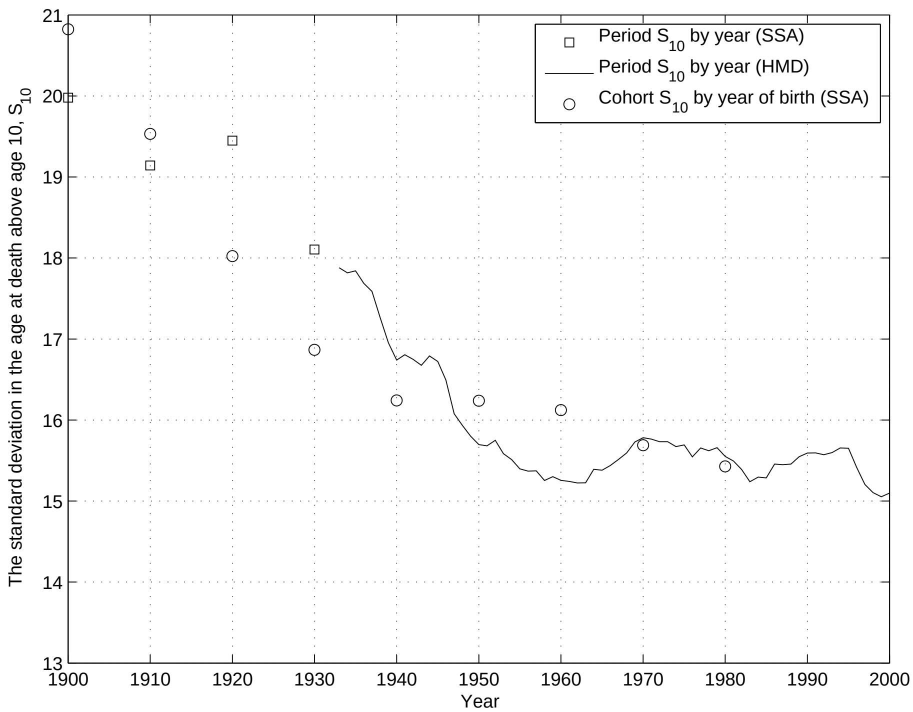

The underlying data are period and cohort life table death distributions, both historical and forecast, taken from Bell and Miller (2005), labeled SSA in the graph, and Human Mortality Database (2009), labeled HMD. The statistics in the figure are standard deviations in length of life above age 10, S_{10}, which are calculated as described by Edwards and Tuljapurkar (2005).

4.3 Population health over time and space

By providing a new method of converting units of variance in length of life into units of mean, equation (14) facilitates a more complete assessment of population health in a variety of settings. In particular, we can reexamine either the differences between advanced countries such as the U.S. and Sweden at a point in time, the differences between observations of the same advanced country at different points in time, or the differences between rich and poor countries across time. Applied intertemporally, this new perspective also provides insights into the nature and timing of demographic and epidemiological transitions that have occurred and are proceeding today in various parts of the world.

4.3.1 Contemporary differences between advanced countries

In 1999, individuals in the U.S. experienced a standard deviation in life spans conditional on survival to age 10 equal to about $\mathrm{S}_{10} = 15$ years. At that level, each year of standard deviation is worth about $p_S = -0.45$ year of mean life span in this model, assuming $r = \delta = 0.03$, the standard value in calibration exercises using U.S. data. In Sweden that same year, S_{10} was about 13. According to this model, 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.[30] 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. If we account for differences in S_{10}, the total difference in population health between the U.S. and Sweden, as measured by "effective" life expectancy, is more like 3.2 life years per person rather than 2.3, an increase of more than a third.

A similar result emerges when we examine a group of advanced economies. Among the 27 members of the OECD that are also designated as high-income countries by the World Bank, life expectancy at birth averaged 78.4 with a standard deviation of about 2 years in 2000, while S_{10} averaged 13.5 with a standard deviation of 0.5 year. Translating the differences in S_{10} from the average into penalties in life expectancy using $p_S = -0.45$ widens the standard deviation in effective life expectancy from 2 to about 2.2 years, or by roughly 10 percent. Measured inequality among advanced countries increases when we account for different variances, but the effect is moderate in size.

---

by which we mean the uncertainty in life span in a known probability distribution, because the time-series evidence seems to support that conclusion (Lee and Carter, 1992). Using the Lee-Carter method of forecasting mortality, I found that forecast uncertainty appears to be small, perhaps 1 year in standard deviation for the cohort born in 2000, relative to life-table uncertainty around 15.3 years. Since these are independent risks, this cohort's total $\mathrm{S}_{10} = 15.33$, or only 0.03 year higher than that implied by the median forecast life table.

[30]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 $M_1 - M_2 = \hat{\delta}(S_1^2 - S_2^2)/2$.

4.3.2 Temporal trends within a country

Because the underlying trends in mean and adult variance are so different, as shown by Edwards and Tuljapurkar (2005), the findings in this paper suggest it is worthwhile to decompose the economic value of historical gains against mortality into portions attributable to reductions in variance versus improvements in mean life span. Nordhaus (2003) and Murphy and Topel (2006) measure the total value of mortality improvement in the U.S. over historical periods using the entire survivorship curve. They find that the value of health improvements is very large, rivaling the value of GDP. But how much is due to increases in the mean, which show no signs of stopping (White, 2002; Oeppen and Vaupel, 2002), and how much is due to declines in adult variance, which have largely stalled since 1960?

The top two rows of Table 1 list e_0 and S_{10} in 1900, 1950, and 2000. Survivorship to age 10, ℓ_{10}, which acts as an importance weight for characteristics of the life table above age 10, is shown in the third row. Changes over time in these three life-table parameters appear in the bottom panel. Between 1900 and 1950, life expectancy rose by 20.7 years from 47.7 to 68.4, an increase of over 40 percent, while S_{10} fell by 4.4 years, from 20 to 15.6, a reduction of about 20 percent. Life expectancy continued to increase steadily after 1950, rising an additional 8.3 years by 2000, but further reductions in S_{10} were slight, totaling only 0.9 year over the second half of the century.

The bottom panel in Table 1 translates the observed declines in S_{10} into equivalent gains in mean life years using the formula for p_S derived earlier, assuming $\delta = 0.03$ in all periods. The relevant price of a discrete change in S_{10} is δ times the average S_{10} during the interval, and the resulting value of the decline in S_{10} must be discounted by the average level of survivorship to age 10 in order to capture the benefit to the average individual. This latter step is especially important when ℓ_{10} is relatively low, as it was in the U.S. early last century. Combining the value of decreased S_{10} with the change in e_0 produces a rough measure of the total gains against mortality in terms of mean life years. By this accounting, reductions in life-span variance prior to 1950 accounted for 9 percent of the gains in effective life expectancy, but after 1950 that figure fell to 4.6 percent.

4.3.3 Convergence in "full income" across countries

In a widely cited paper, Becker, Philipson and Soares (2005) reveal relatively more global convergence in human well-being when it is measured by "full income," a statistic they devise that combines GDP per capita with the monetized value of life, than when it is measured by GDP per capita alone. Owing to data constraints in their panel of 96 countries measured since 1960, they can only account for changes in life expectancy at birth rather than changes in higher moments or in entire survivorship schedules. Given the costliness of variance, has its omission biased their results, and if so which way?

Complete life tables are not widely available for a broad cross section of

Table 1: Changes in U.S. life expectancy, adult life-span variance, and survivorship since 1900

		1900	1950	2000
	Life expectancy, e_0	47.7	68.4	76.7
	Std. dev. in adult life span, S_{10}	20.0	15.6	14.7
	Survivorship to age 10, ℓ_{10}	0.782	0.963	0.991

		1900-2000	1900-1950	1950-2000
[1]	Average S_{10}	17.4	17.8	15.2
[2]	Average p_S	0.52	0.53	0.45
[3]	Change in S_{10}, ΔS_{10}	5.3	4.4	0.9
[4]	Life year benefit of ΔS_{10}, [2] $\times$ [3]	2.8	2.3	0.4
[5]	Average ℓ_{10}	0.886	0.872	0.977
[6]	Life year benefit wtd. by avg. ℓ_{10}, [4] $\times$ [5]	2.4	2.0	0.4
[7]	Change in life expectancy, Δe_0	29.0	20.7	8.3
[8]	Total improvement in life years, [6] + [7]	31.4	22.7	8.7
[9]	Share due to ΔS_{10}, [6] $\div$ [8]	0.078	0.090	0.046

Demographic data are simple averages of sex-specific period life tables presented by Bell and Miller (2005) and are based on age-specific mortality rates measured in the given year. Life expectancy at birth, e_0, is the familiar average number of years lived starting from birth or age 0. The standard deviation of length of life conditional on survival to age 10, S_{10}, is calculated using rescaled probabilities of death above age 10 and is measured around the mean length of life conditional on survival to age 10, M_{10}, which equals $e_{10}+10$. The survivorship probability at age 10 is ℓ_{10}. The price of S_{10} in terms of mean life years is given by $p_S = -\delta S_{10}$, as described in the text, where it is assumed that $\delta = 0.03$ in all periods.

countries because the underlying mortality data are typically of poor quality. For many developing countries where data is scarce, surveillance agencies have estimated levels of life expectancy at birth, e_0, using model life tables, but historical databases do not include complete model life tables. Recently, Lopez et al. (2002) and the World Bank have produced estimates of current life tables for virtually all countries, but historical coverage has remained lacking. New work by Edwards (2011) collects and reconstructs historical estimates of life tables in developing countries and allows a reassessment of convergence in the full income measure after accounting for changes in variance and the shape of the survivorship function.[31] As discussed by Edwards, these new data are of varying quality, but they are no worse than much of the original data on life expectancy commonly found in databases and used by Becker, Philipson and Soares (2005) and many others.

The methodological component in Becker, Philipson and Soares (2005) that captures the effects of mortality is the value of an annuity based on the survivorship function, $\ell(t)$, given by

$$A(\ell) = \int_0^\infty e^{-rt}\ell(t)\ dt. \tag{17}$$

With full annuitization and the time discount rate δ equal to the interest rate r, indirect utility is the product of $A(\ell)$ and the period utility function. When only life expectancy e_0 is known, survivorship must implicitly be rectangular, or equal to 1 at all ages until falling to 0 at e_0. Under those conditions, equation (17) simplifies to

$$A^*(\ell) \approx \frac{1 - e^{-r\cdot e_0}}{r}. \tag{18}$$

As seen and discussed earlier, for example in equation (11), the presence of variance in adult length of life will reduce the value of the annuity:

$$A^\dagger(\ell) \approx \frac{1 - e^{-r\cdot e_0 + r^2 {S_{10}}^2/2}}{r}, \tag{19}$$

where S_{10} is the standard deviation in adult length of life. If there is any adult variance, $A(\ell)$ and $A^\dagger(\ell)$ are both less than $A^*(\ell)$, so equation (19) should be a closer approximation to (17) than (18).

As discussed in section 3, high infant mortality also tends to reduce $A(\ell)$, holding other things including e_0 equal. Indeed, infant mortality is additional variance in human life span that is highly non-Gaussian and non-central. In the simple model based on Yaari (1965), living with complete certainty from birth until death at age e_0 is more valuable than facing a nonzero probability d of dying in infancy followed by living with certainty until $e_0/(1-d) > e_0$,

[31] In their original work, Becker, Philipson and Soares (2003) examined a narrower cross section of 49 countries in the WHO mortality database. With data on age-specific mortality rates, they constructed life tables and full survivorship schedules. Becker, Philipson and Soares (2005) measured convergence among a broader sample of 96 countries for whom only life expectancy at birth is provided by the World Bank's World Development Indicators database.

even though life expectancy at birth is exactly the same in either case. This is because the cost of infant mortality outweighs the benefit of living past e_0 due to discounting, the same reason why S_{10} is costly. Both equations (18) and (19) will typically overestimate $A(\ell)$ because each formula omits infant mortality, which has been and remains high in the developing world. If complete life tables are available, of course, it is possible to avoid this problem by measuring $A(\ell)$ exactly. In order to gauge the importance of life-span uncertainty for convergence, I estimate $A(\ell)$ and full income using all three methods and then compare results.

Tables 2 and 3 reproduce their counterparts in Becker, Philipson and Soares (2005) using a wider sample of 180 countries covering virtually all of the world's population. Results using only the original 96 countries examined by Becker, Philipson and Soares are similar and indicate slightly more convergence than is shown here.[32] The top panel in Table 2 shows population weighted averages of aggregates from life tables and national income accounts for a set of world regions as defined by the World Bank, for the poorest and richest countries in 1970, and for the world as a whole, using mortality data from Edwards (2011) and income data from the Penn World Table, Maddison (2003), and the IMF. All regions gained income per capita during the interval, at an average annual rate of about 1.8 percent worldwide, and all but one gained in terms of life expectancy, which rose 0.27 year per annum on average. The world standard deviation in adult life span remained fairly steady, declining only 0.2 year from 17.0 to 16.8. Survivorship to age 10 rose 7 percentage points, from 86.7 to 93.7 percent. Although almost all regions experienced gains against mortality, there was also clear heterogeneity. Europe and Central Asia, a World Bank category that does not include the high income European countries, experienced only a very small 0.9 year total gain in life expectancy and actually suffered a small

[32] In general, however, convergence results are not robust to the breadth and quality of the sample, nor to the use of population weights, and they appear not to have been robust in the original data examined by Becker, Philipson and Soares (2005) either. Unweighted estimates, which place much more emphasis on the experiences of small countries, reveal greater inequality in full income than in GDP per capita in 2000. This is true in the broad sample of 180, in the smaller sample of 96 used by Becker, Philipson and Soares, and it is also true in a subset of 35 countries with high-quality mortality data drawn from the Human Mortality Database (2009). In the last subsample, which is dominated by high-income countries, even population-weighted estimates show slightly increased inequality in full income compared to GDP per capita in 2000. Two factors probably account for these patterns. First, countries that have been hit hardest by HIV/AIDS tend to be small. Unweighted estimates are therefore likely to overestimate the impact of the disease on convergence in average human well-being, although they correctly measure the convergence across countries. Second, because high-quality mortality data is a luxury affordable only to high-income countries, and income also correlates strongly with the demographic transition, it is not surprising that mortality trends based only on high-quality data are different than those based on a much broader sample. Still, it is troubling that the basic result, namely increased convergence in full income compared to GDP per capita, seems to hinge on the use of mortality data of questionable quality. Recent developments in a related literature on income inequality and population health testify to the inherent dangers (Judge, Mulligan and Benzeval, 1998). Less troubling but still worthy of note is the finding that average well-being in small countries seems not to be converging toward that in large countries, which bears broad implications, such as for the continued spread of the demographic transition.

increase in S_{10}, from 15.9 to 16.2, although ℓ_{10} increased somewhat. Sub-Saharan Africa gained almost 5 years in life expectancy, due to a large increase in ℓ_{10}, which rose from 74.1 to 82.7 percent. But variance in adult life dropped only 0.4 to 19.4.

The bottom panel of Table 2 translates gains in income and survivorship into gains in the "full income" measure proposed by Becker, Philipson and Soares (2005) using the three different methods of valuing life span discussed above. When only e_0 is used, the annual value of gains is an additional \$800 (in 2000 international dollars) for the world as a whole, and the growth in full income is 2.2 percent, or 0.4 faster than the growth rate of income per capita during the period. Accounting for both e_0 and S_{10} increases the contributions of mortality gains, to \$1,159 and 0.5 percentage point faster annual growth respectively. Using the entire survivorship schedule including infant mortality more than doubles the gains, to \$1,684 and an extra 0.7 percentage point in annual growth.

The hefty importance these comparisons seem to attach to S_{10} is odd in light of the relatively small decrease in world S_{10} since 1970, and it is in fact misleading. The presence of any S_{10} at all raises the value of gains in e_0 by increasing their marginal utility. Equation (15) shows this clearly: the marginal utility of the mean length of life, which appears in the denominator, is a decreasing function of the mean but an increasing function of the variance. When I account for S_{10} in addition to e_0 in Table 2, I am basically turning variance on after being off, which has a large effect on measured trends in well-being even when the trend in variance is negligible.

A more appropriate way to decompose growth in full income into portions attributable to changes in individual moments is by measuring gains using counterfactual but fully realistic survivorship curves that hold one or more moments constant over time. Without any change in world S_{10}, the value of total annual survivorship gains would be about 5 percent lower than the \$1,685 we observe in Table 2. This share is similar to what I found in Table 1 for the U.S. after 1950. Indeed, average annual declines in S_{10} for the U.S. since 1950 and for the world since 1970 are both only around 0.01 to 0.02, a sobering remark on recent progress against high adult variance in the developing world. Declines in infant and child mortality, on the other hand, have been large and more prevalent among poor countries (Moser, Shkolnikov and Leon, 2005). To gauge their importance, I generated a fictitious world survivorship curve for 2000 assuming there was change since 1970 only in ℓ_{10}, and not in the shape of the life-span distribution above age 10. Based on this counterfactual, the value of improvements spurred by increases in ℓ_{10} accounted for more than 75 percent of the total value of gains against mortality.

Heterogeneity across regions in the progress against mortality is certainly evident in the lower panel in Table 2, where results broadly amplify the patterns reported by Becker, Philipson and Soares (2005). In dollar terms, the benefits of mortality reduction were felt considerably more strongly by the richest 50 percent of countries, but growth rates of full income were faster for the poorest 50 percent because their initial money incomes were so much lower. Sub-Saharan

Table 2: Value of survivorship gains by region of the world and groups of countries, 1970–2000

	1970				2000			
	Life exp. e_0	Life std. dev. S_{10}	Survi-vorship ℓ_{10}	GDP per capita	Life exp. e_0	Life std. dev. S_{10}	Survi-vorship ℓ_{10}	GDP per capita
East Asia & Pacific	58.3	16.4	0.871	695	69.7	15.1	0.959	3,908
Europe & Central Asia	67.3	15.9	0.947	5,498	68.2	16.2	0.970	6,838
High income	70.6	15.0	0.971	12,951	77.7	14.4	0.992	25,954
Latin America & Caribbean	60.4	16.8	0.881	4,839	71.5	16.7	0.965	7,085
Middle East & North Africa	53.7	17.9	0.822	3,202	67.7	15.2	0.947	4,569
South Asia	47.8	17.7	0.757	1,183	60.9	17.1	0.898	2,510
Sub-Saharan Africa	45.8	19.8	0.741	1,565	50.7	19.4	0.827	1,667
Poorest 50% countries in 1970	53.6	17.4	0.820	1,098	64.1	16.8	0.919	3,326
Richest 50% countries in 1970	67.9	15.6	0.948	10,105	73.4	16.0	0.978	17,024
World	58.8	17.0	0.867	4,360	66.9	16.8	0.937	7,505

	Value of survivorship gains in annual income calculated with:			Yearly growth rate of full income (%) calculated with:		
	only e_0	e_0 and S_{10}	entire $\ell(x)$	only e_0	e_0 and S_{10}	entire $\ell(x)$
East Asia & Pacific	477	720	987	6.3	6.5	6.7
Europe & Central Asia	60	54	281	0.8	0.8	0.9
High income	2,035	2,766	3,343	2.6	2.7	2.8
Latin America & Caribbean	926	1,299	2,009	1.7	1.8	2.1
Middle East & North Africa	838	1,474	1,903	1.8	2.1	2.4
South Asia	438	699	877	3.1	3.4	3.6
Sub-Saharan Africa	105	197	235	0.4	0.6	0.7
Poorest 50% countries in 1970	425	651	865	4.2	4.4	4.6
Richest 50% countries in 1970	1,067	1,290	2,020	2.0	2.0	2.1
World	800	1,159	1,684	2.2	2.3	2.5

Notes: The sample, fully described in Edwards (2011), comprises 180 countries including Taiwan. Regions are otherwise as defined by the World Bank, and regional averages are weighted by population. GDP per capita is measured in 2000 international prices, adjusted for terms of trade. Measures are collected from the Penn World Table, Maddison (2003), and the IMF. Life expectancy at birth, e_0, the standard deviation of length of life above age 10, S_{10}, survivorship to age 10, ℓ_{10}, and the entire $\ell(x)$ distributions are derived from official data and estimates collected by Edwards (2011). The value of life expectancy gains and full income are calculated three ways using the methodology and parameter values of Becker, Philipson and Soares (2005) and 1970 as the base year. When only e_0 is used, survivorship is assumed to be 1 until dropping to zero at age e_0. When e_0 and S_{10} are used, deaths are assumed to be distributed normally around a mean of e_0. When the entire $\ell(x)$ curve is used, the survivorship schedule is fully realistic and reflects everything: the mean, the spike in deaths due to infant and child mortality, the old-age hump, and all its skewness and kurtosis.

Table 3: Evolution of cross-country inequality in full income, 1970–2000

	Income per capita		Full income in 2000 calculated with:		
	1970	2000	e_0 only	e_0 and S_{10}	entire $\ell(x)$
Relative mean deviation	0.4816	0.4227	0.4117	0.4038	0.3970
Coefficient of variation	1.2041	1.1801	1.1524	1.1342	1.1076
Std. dev. of logs	1.2216	0.9910	0.9874	0.9791	0.9819
Gini coefficient	0.5439	0.4999	0.4904	0.4836	0.4781
Regression to the mean since 1970		−0.3507	−0.3681	−0.3805	−0.3872

Notes: See notes to Table 2. Inequality measures are weighted by country population. Regression to the mean is the coefficient from a weighted OLS regression of the change in the natural log of income over the period on its initial logged level, with 1970 populations as weights. All four regression coefficients are statistically significant at the 1% level.

Africa remains a basket case even when improvements in the entire survivorship curve are priced, raising growth in full income to only 0.7 percent. But a similar picture emerges for Europe and Central Asia, where the negative influence of a slight increase in S_{10} was more than offset by improvements in ℓ_{10}, but not by enough to raise growth past 0.9 percent. Elsewhere, economic growth that was already fairly robust is further enhanced by accounting for improvements in survivorship. In East Asia and the Pacific, growth in full income is a stout 6.7 percent after accounting for much improvement in all three mortality statistics, up 0.4 from the e_0-only reading. The Middle East, North Africa, and South Asia benefit the most, about 0.6 percentage point in annual growth, from accounting for all changes in $\ell(x)$ as opposed to e_0 alone.

Table 3 descends past regions to examine inequality in average human well-being using countries weighted by population as the unit of observation. By any one of several measures, inequality across individuals in full income in 2000 is less than inequality in income per capita in 2000 or in 1970. The Gini coefficient, a widely cited statistic, registers 0.4781 when measured using full income in 2000 derived with countries' entire $\ell(x)$ schedules. By comparison, the Gini on per capita income in 2000 was 0.4999. Accounting for successively higher moments of the life-span distribution reinforces this basic convergence result, as it did among regions in Table 2. Another upshot of the math in equation (15) is that the value of a given increase in mean life span will be larger when S_{10} is higher, such as in a poor country. Part of what we see here is thus surely the mechanical effect of S_{10} and infant mortality raising the measured value of gains in e_0 disproportionately more among poor countries with high variance. But surely also at work are continued global convergence in infant mortality even as convergence in e_0 turned to divergence after 1980 (Moser, Shkolnikov and Leon, 2005), as well as some convergence in country S_{10} revealed by Edwards (2011) as narrowing within-country inequality.

These results mirror and amplify those of Becker, Philipson and Soares (2005) on convergence in average human well-being, but some circumspection is in order. Considerably less progress against mortality in poor but more sparsely populated areas like Sub-Saharan Africa, parts of Europe, and Central Asia, was not sufficient to reduce overall convergence in average human full income due to considerable progress in larger countries. Still, population-weighted estimates suggest more convergence in average human well-being since 1970 than in average income alone. Accounting for gains against life-span variance, either in the form of S_{10} or infant mortality, strengthens the findings of Becker, Philipson and Soares for two reasons. Mechanically speaking, correctly accounting for even a stable level of variance raises the value of increases in life expectancy, and that effect increases with the level of variance, which is higher in poor countries. In addition, total variance was also strongly declining over this period, primarily due to robust declines in infant mortality but also due to moderate declines in S_{10}. The gains against variance, broadly defined, contributed a large amount to the value of progress against mortality.

4.3.4 A new perspective on the demographic transition

The insights of this paper also suggest a new interpretation of the historical timing of age-specific gains against disease and mortality during the demographic and epidemiologic transitions. As summarized by Wilmoth (2003), the classic transition begins with a decline in infant mortality and early death, brought about by progress against infectious disease. That is, the first stage of progress drastically lessens the unconditional variance in life spans. The second stage of the transition is characterized by a shift in focus away from infectious disease and toward treating chronic degenerative diseases afflicting the elderly. This works to lengthen the average adult life span but probably does not reduce the variance much if at all.

The current framework suggests this sequence was probably optimal, if it were not practically required.[33] High levels of variance inflate its cost considerably. Combating infectious disease reduces unconditional variance directly by reducing infant and child mortality, and it also reduces uncertainty in adult life. Certainly in the context of the modern developing world, in which life expectancy is short, variance is high, both infectious and chronic diseases claim lives, and technologies and practices to combat either one are readily available via the developed world, revealed preference seems to indicate that public health priorities largely center on the unconditional variance first, i.e., infant and premature adult mortality. An outlier based on this perspective, HIV/AIDS is a hybrid disease afflicting old and young alike.

[33] A recent paper posits that contemporary declines in adult mortality are due to earlier declines in childhood disease and mortality (Finch and Crimmins, 2004), owing to a reduced accumulation of inflammatory exposure from infectious disease. One might infer that sustained mortality decline could happen in no other way. But in many modern developing countries, infant mortality remains high even though imported medical technologies can reduce adult mortality. Life expectancy in many such countries has increased (Becker, Philipson and Soares, 2005), although perhaps not as robustly as in the developed world.

Historically, these two very different stages of the transition produced a seamless pattern of steady increases in life expectancy at birth over time (Oeppen and Vaupel, 2002), which fits well with steady growth in per capita incomes (Hall and Jones, 2007) and shows that average health outcomes were rising consistently. But the technologies, cost structures, and incidence of benefits during each phase were entirely different. How societies set priorities in achieving mortality decline is a major question. A key insight of this paper, that variance in life span is costlier relative to mean life span when variance is higher, suggests that declines in total variance should precede sustained progress in the adult mean or mode.[34]

4.4 Uncertain life span and economic behavior

This paper, like others in the literature on the value of gains against mortality, explores only one economic perspective on the cost of uncertainty in life span: the willingness to pay. Although willingness to pay should in principle encompass or account for all expected behavioral responses to changes in uncertainty, my simple theoretical model only captures the response of consumption and saving and not other behavioral responses that we think might be very interesting. The latter are a burgeoning field of research. Kalemli-Ozcan and Weil (2010) examine the effects of uncertain life span in a model with endogenous labor supply. Their results suggest that an optimal policy under high uncertainty might be low saving and working until death, or what macroeconomists might consider extreme rule-of-thumb as opposed to life-cycle behavior. Patterns in retirement behavior early in the 20th century may support this view (Costa, 1998). But in modern data, Hurd, Smith and Zissimopoulos (2004) find small effects of low subjective survivorship probabilities on retirement behavior among older Americans in the Health and Retirement Study. Education, another form of saving, is also endogenous and may react to uncertainty in length of life. Building on earlier work by Kalemli-Ozcan, Ryder and Weil (2000) that focuses primarily on the average length of life, Li and Tuljapurkar (2004) show that uncertainty has an effect on educational attainment in a general equilibrium setting. These and other directions are promising avenues for further research.

5 Conclusion

In the standard model of time-separable utility introduced by Yaari (1965), uncertainty in life span is costly when the force of time discounting, δ, which is also approximately the coefficient of absolute risk aversion in life span, is positive and sufficiently near the real interest rate. Even when wealth is fully annuitized, individuals with these preferences are hurt by uncertainty in life span and would

[34] To be sure, the story must also involve marginal costs, since the socially optimal allocation of resources occurs where the ratio of marginal utilities is equal to the ratio of marginal costs. Reducing variance through improved sanitation or other public health initiatives is likely to be much less costly than increasing the adult mean, which involves treating degenerative diseases.

be willing to trade away $p_S = -\delta S$ years of mean life span in return for one less year in standard deviation, S. Because average life expectancy is increasing linearly over time in advanced countries while the standard deviation in adult life span now roughly fixed (Edwards and Tuljapurkar, 2005), constant absolute risk aversion is consistent with stable risk premia and thus is an intuitive result.[35] Risk aversion over length of life also fits the relatively scant empirical evidence on stated preferences and behavior in medical settings. If $\delta = r = 0.03$, which is a standard parameterization, the average American would be willing to give up 0.45 life year in return for one year less in standard deviation, which is currently about 15 when measured by S_{10}, the standard deviation in length of life above age 10. This is large, implying that differences in population health between the U.S. and Sweden are more like 3.2 life years, or 40 percent higher than the difference of 2.3 years we find in life expectancy alone.

Adult variance is interesting because of the large differences in levels and trends we see among industrialized countries today (Edwards and Tuljapurkar, 2005), but it is not the only source of variance. Infant mortality is low in advanced countries today, but it was very high during historical periods and is high in many modern developing countries. Discounting places a very high cost on infant mortality indeed, making it two to three times more costly than the reduction in mean length of life it represents. This is because all of the probability weight at play occurs so early in life, during years that are extremely valuable in the simple model of Yaari (1965). This result would be moderated if the empirical age trajectory of the value of a statistical life year were superimposed, because it is hump-shaped (Aldy and Viscusi, 2008). But without accounting for the unconditional variance in length of life represented by infant deaths, valuing gains in life expectancy at birth will significantly understate the benefit of declining infant mortality.

The framework developed in this paper allows me to decompose the overall value of progress against mortality into parts attributable to reductions in infant mortality, increases in the adult mean, and reductions in adult variance. This decomposition is interesting because gains against adult variance and infant mortality have largely stopped in developed countries (Edwards and Tuljapurkar, 2005), while increases in the mean continue apace. Worldwide, moderate declines in adult variance combined with broad-based reductions in infant mortality, the larger source of unconditional variance in human life span,

[35]The intuition derives from an analogy to the relationships between time trends in aggregate consumption, risk premia in consumption, and preferences over consumption, which are described by Campbell and Viceira (2002). Log consumption is trending linearly upward over time, while financial risk premia and the standard deviations of asset returns and log consumption have remained roughly constant. Preferences over consumption that exhibit constant relative risk aversion are consistent with these facts, and this is one reason why the familiar power utility function is useful for modeling. By comparison, average length of life is trending upward at an approximately linear rate, and the standard deviation has remained roughly constant since 1960. These facts imply that constant absolute risk aversion over length of life would be consistent with these facts and with stable risk premia placed on length of life. Although we technically do not observe the latter in any market, it seems reasonable to expect that they probably are constant.

have produced much global convergence in average human well-being, more than is implied by trends in e_0 alone. These findings amplify those in the pioneering work of Becker, Philipson and Soares (2005), who due to data limitations could only consider life expectancy. If the insights in this paper were applied to the broader definition of welfare adopted by Jones and Klenow (2010), who also measure only e_0, it is possible their evidence of divergence among countries since 1980 may be partially reversed.

Several limitations of my approach are worth highlighting. I do not account for education or for physical capital, both of which are considered in a general equilibrium setting by Li and Tuljapurkar (2004). I also do not consider variation in morbidity or the quality of life. Bequests are a potentially key omission, because they could significantly reduce the marginal disutility of life-span uncertainty. But even if bequests were intended, which is unclear, they would also reduce the marginal utility of mean life span, leaving an ambiguous effect on the price of life-span variance relative to the mean. I also make no allowance for the special psychology that we know surrounds death (Slemrod, 2003). One could argue that knowing the precise date of death is actually not a good, at least in terms of emotional benefit, and that some uncertainty is preferable. Still, the economic cost of uncertain length of life in terms of retirement and estate planning is real and plausibly much larger. Questions remain about time discounting and its relationship to risk preferences over periods of life. Bommier (2006, 2008) and Bommier and Villeneuve (2011) develop a more general model than Yaari's and report a greater degree of curvature than implied by exponential time discounting alone. In their study of 30 women in perfect health asked to rank lotteries over life span, Verhoef, Haan and van Daal (1994) report evidence supporting prospect theory: risk-seeking behavior over small gambles and risk aversion over large. My estimate of the cost of uncertain life span is probably a conservative estimate, and by presenting it I hope to motivate further research into this topic.

The primary implication for policy is that uncertainty in length of life is costly and should be targeted when it is high, whether in entire countries or among specific subgroups. Because its cost rises with its level, and since the level of uncertainty also diminishes the value of gains in life expectancy, policies that reduce uncertainty should be favored over those that increase average length of adult life when uncertainty is high, such as when infant mortality is high. A more provocative interpretation of this bottom line is that a high-variance country like the U.S. may stand to gain by focusing more on spreading the existing benefits of health treatments and technology more broadly across its citizens, at the expense of investing in discovery of new treatments and technology that extend life. Given the key role attached to U.S. pharmaceutical innovation in driving old-age mortality decline worldwide (Lichtenberg, 2007), such a refocusing may not come without cost. A less provocative implication is that developing countries are better served by attacking the causes of high uncertainty first, infectious diseases and infant mortality, before extending assistance that extends life for adults. In large part, development assistance is already configured to do this.

References

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

Aldy, Joseph E. and W. Kip Viscusi. 2008. "Adjusting the Value of a Statistical Life for Age and Cohort Effects." *Review of Economics and Statistics* 90(3):573–581.

Atkinson, Anthony B. 1970. "On the Measurement of Inequality." *Journal of Economic Theory* 2(3):244–263.

Banks, James, Richard Blundell and Sarah Tanner. 1998. "Is There a Retirement Savings Puzzle?" *American Economic Review* .

Barsky, Robert B., F. Thomas Juster, Miles S. Kimball and Matthew D. Shapiro. 1997. "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study." *Quarterly Journal of Economics* pp. 537–579.

Becker, Gary S. and Casey B. Mulligan. 1997. "The Endogenous Determination of Time Preference." *Quarterly Journal of Economics* 112(3):729–758.

Becker, Gary S., Tomas J. Philipson and Rodrigo R. Soares. 2003. "The Quantity and Quality of Life and the Evolution of World Inequality." *NBER Working Paper* 9765.

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.

Bernheim, B. Douglas, Andrei Shleifer and Lawrence H. Summers. 1985. "The Strategic Bequest Motive." *Journal of Political Economy* 93(6):1045–1076.

Bishai, David M. 2004. "Does time preference change with age?" *Journal of Population Economics* 17(4):583–602.

Bleichrodt, Han and John Quiggin. 1999. "Life-cycle preferences over consumption and health: When is cost-effectiveness analysis equivalent to cost-benefit analysis?" *Journal of Health Economics* 18(6):681–708.

Bleichrodt, Han and Magnus Johannesson. 1997. "The Validity of QALYs: An Experimental Test of Constant Proportional Tradeoff and Utility Independence." *Medical Decision Making* 17(1):21–32.

Bommier, Antoine. 2006. “Uncertain Lifetime and Intertemporal Choice: Risk Aversion as a Rationale for Time Discounting.” *International Economic Review* 47(4):1223–1246.

Bommier, Antoine. 2008. “Rational Impatience?” *Unpublished working paper* .

Bommier, Antoine and Bertrand Villeneuve. 2011. “Risk Aversion and the Value of Risk to Life.” *The Journal of Risk and Insurance,* forthcoming.

Campbell, John Y. and Luis M. Viceira. 2002. *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*. Oxford: Oxford University Press.

Chang, Fwu-Ranq. 1991. “Uncertain Lifetimes, Retirement and Economic Welfare.” *Economica* 58(230):215–232.

Costa, Dora L. 1998. *The Evolution of Retirement: An American Economic History, 1880–1990*. Chicago: University of Chicago Press.

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, Angus Deaton and Adriana Lleras-Muney. 2006. “The Determinants of Mortality.” *Journal of Economic Perspectives* 20(3):97–120.

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

Deaton, Angus. 2004. Health in an Age of Globalization. In *Brookings Trade Forum 2004: Globalization, Poverty, and Inequality*, ed. Susan M. Collins. Chicago: University of Chicago Press pp. 83–110.

Dynan, Karen E., Jonathan Skinner and Stephen P. Zeldes. 2004. “Do the Rich Save More?” *Journal of Political Economy* 112(2):397–444.

Edwards, Ryan D. 2009. “The Cost of Cyclical Mortality.” *B.E. Journal of Macroeconomics,* 9(1):Contributions, Article 7.

Edwards, Ryan D. 2011. “Changes in World Inequality in Length of Life: 1970–2000.” *Population and Development Review,* 37(3):499–528.

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.

Ehrlich, Isaac. 2000. “Uncertain Lifetime, Life Protection, and the Value of Life Saving.” *Journal of Health Economics* 19:341–367.

Epstein, Larry G. and Stanley E. Zin. 1989. “Substitution, Risk Aversion, and the Temporal Behavior of Consumption Growth and Asset Returns I: A Theoretical Framework.” *Econometrica* 57(4):937–969.

Epstein, Larry G. and Stanley E. Zin. 1991. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis." *Journal of Political Economy* 99(2):263–286.

Finch, Caleb E. and Eileen M. Crimmins. 2004. "Inflammatory Exposure and Historical Changes in Human Life-Spans." *Science* 305(5691):1736–1739.

Fuchs, Victor R. 1982. Time Preference and Health: An Exploratory Study. In *Economic Aspects of Health*, ed. Victor R. Fuchs. Chicago: University of Chicago Press. Chapter 3.

Fuchs, Victor R. and Hal Ersner-Hershfield. 2008. "Inequality in Length of Life." *SIEPR Policy Brief* October 1.

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

Hall, Robert E. and Charles I. Jones. 2007. "The Value of Life and the Rise in Health Spending." *Quarterly Journal of Economics* 122(1):39–72.

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.

Hubbard, R. Glenn and Kenneth L. Judd. 1987. "Social Security and Individual Welfare: Precautionary Saving, Borrowing Constraints, and the Payroll Tax." *American Economic Review* 77(4):630–646.

Human Mortality Database. 2009. 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.

Hurd, Michael D., James P. Smith and Julie Zissimopoulos. 2004. "The Effects of Subjective Survival on Retirement and Social Security Claiming." *Journal of Applied Econometrics* 19(6):761–775.

Ibbotson Associates. 2002. *Stocks, Bonds, Bills, and Inflation Yearbook.* Chicago: Ibbotson Associates.

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.

Jones, Charles I. and Peter J. Klenow. 2010. "Beyond GDP? Welfare Across Countries and Time." *NBER Working Paper* 16352.

Judge, Ken, Jo-ann Mulligan and Michaela Benzeval. 1998. "Income Inequality and Population Health." *Social Science & Medicine* 46(4–5):567–579.

Kalemli-Ozcan, Sebnem and David N. Weil. 2010. "Mortality Change, the Uncertainty Effect, and Retirement." *Journal of Economic Growth* 15(1):65–91.

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):1–23.

Kelley, Amy S. 2011. "Treatment intensity at end of life — time to act on the evidence." *Lancet* 378(9800):1364–1365.

Kitagawa, Evelyn M. and Philip M. Hauser. 1973. *Differential Mortality in the United States: A Study in Socioeconomic Epidemiology.* Cambridge, MA: Harvard University Press.

Kopczuk, Wojciech and Joseph P. Lupton. 2007. "To Leave or Not To Leave: The Distribution of Bequest Motives." *Review of Economic Studies* 74(1):207–235.

Kreps, David M. and Evan L. Porteus. 1978. "Temporal Resolution of Uncertainty and Dynamic Choice Theory." *Econometrica* 46(1):185–200.

Kwok, Alvin C, Marcus E Semel, Stuart R Lipsitz, Angela M Bader, Amber E Barnato, Atul A Gawande and Ashish K Jha. 2011. "The intensity and variation of surgical care at the end of life: a retrospective cohort study." *Lancet* 378(9800):1408–1413.

Lee, Ronald D. 2003. "The Demographic Transition: Three Centuries of Fundamental Change." *Journal of Economic Perspectives* 17(4):167–190.

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.

Lichtenberg, Frank R. 2007. "The Impact of New Drugs on US Longevity and Medical Expenditure, 1990–2003: Evidence from Longitudinal, Disease-Level Data." *American Economic Review* 97(2):438–443.

Lopez, Alan D., O. B. Ahmad, Michel Guillot, M. Inoue, B. D. Ferguson, J. A. Salomon, Christopher J. L. Murray and Kenneth H. Hill. 2002. *World Mortality in 2000: Life Tables for 191 Countries.* Geneva: World Health Organization.

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

Maas, Arne. and Peter. P. Wakker. 1994. "Additive conjoint measurement for multiattribute utility." *Journal of Mathematical Psychology* 38(1):86–101.

Maddison, Angus. 2003. *The World Economy: Historical Statistics*. Paris: OECD Development Cenre.

McNeil, Barbara. J., Ralph Weichselbaum and Stephen G. Pauker. 1978. "Fallacy of the five-year survival in lung cancer." *New England Journal of Medicine* 299(25):1397–1401.

Meara, Ellen R., Seth Richards and David M. Cutler. 2008. "The Gap Gets Bigger: Changes in Mortality and Life Expectancy, By Education, 1981–2000." *Health Affairs* 27(2):350–360.

Miyamoto, John. M. and Stephen A. Eraker. 1985. "Parameter Estimates for a QALY Utility Model." *Medical Decision Making* 5(2):191–213.

Moser, Kath, Vladimir Shkolnikov and David A. Leon. 2005. "World mortality 1950–2000: Divergence replaces convergence from the late 1980s." *Bulletin of the World Health Organization* 83(3):202–209.

Murphy, Kevin M. and Robert H. Topel. 2006. "The Value of Health and Longevity." *Journal of Political Economy* 114(5):871–904.

Nordhaus, William D. 2003. The Health of Nations: The Contribution of Improved Health to Living Standards. In *Measuring the Gains from Medical Research: An Economic Approach*, ed. Kevin M. Murphy and Robert H. Topel. Chicago: University of Chicago Press.

Oeppen, Jim and James W. Vaupel. 2002. "Broken Limits to Life Expectancy." *Science* 296:1029–1031.

Picone, Gabriel, Frank Sloan and Donald Taylor Jr. 2004. "Effects of Risk and Time Preference and Expected Longevity on Demand for Medical Tests." *Journal of Risk and Uncertainty* 28(1):39–53.

Pliskin, Joseph S., Donald S. Shepard and Milton C. Weinstein. 1980. "Utility Functions for Life Years and Health Status." *Operations Research* 28(1):206–224.

Ried, Walter. 1998. "QALYs versus HYEs — what's right and what's wrong. A review of the controversy." *Journal of Health Economics* 17(5):607–625.

Rogers, Alan R. 1994. "Evolution of Time Preference by Natural Selection." *American Economic Review* 84(3):460–481.

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

Satchell, Stephen and Susan Thorp. 2011. "Uncertain survival and time discounting: intertemporal consumption plans for family trusts." *Journal of Population Economics* 24(1):239–266.

Shkolnikov, Vladimir M., Evgueni E. Andreev and Alexander Z. Begun. 2003. "Gini coefficient as a life table function: Computation from discrete data, decomposition of differences and empirical examples." *Demographic Research* 8(11).

Slemrod, Joel. 2003. "Thanatology and Economics: The Behavioral Economics of Death." *American Economic Review* 93(2):371–375.

Smits, Jeroen and Christiaan Monden. 2009. "Length of life inequality around the globe." *Social Science & Medicine* 68(6):1114–1123. doi:10.1016/j.socscimed.2008.12.034.

Stiggelbout, A. M., G. M. Kiebert, J. Kievit, J. W. H. Leer, G Stoter and J. C. J. M. de Haes. 1994. "Utility Assessment in Cancer Patients: Adjustment of Time Tradeoff Scores for the Utility of Life Years and Comparison with Standard Gamble Scores." *Medical Decision Making* 14(1):82–87.

Strotz, Robert H. 1956. "Myopia and Inconsistency in Dynamic Utility Maximization." *Review of Economic Studies* 23(3):165–180.

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

Verhoef, Lia C. G., Anton F. J. De Haan and Willem A. J. van Daal. 1994. "Risk Attitude in Gambles with Years of Life: Empirical Support for Prospect Theory." *Medical Decision Making* 14(2):194–200.

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.

White, Kevin M. 2002. "Longevity Advances in High-Income Countries." *Population and Development Review* 28(1):59–76.

Wilmoth, John R. 2003. Human Longevity in Historical Perspective. In *Physiological Basis of Aging and Geriatrics*, ed. Paola S. Timiras. New York: CRC Press pp. 11–24.

Wilmoth, John R. and Shiro Horiuchi. 1999. "Rectangularization Revisited: Variability of Age at Death within Human Populations." *Demography* 36(4):475–495.

Wilson, Chris. 2001. “On the Scale of Global Demographic Convergence 1950–2000.” *Population and Development Review* 27(1):155–171.

Yaari, Menahem E. 1965. “Uncertain Lifetime, Life Insurance, and the Theory of the Consumer.” *Review of Economic Studies* 32(2):137–150.