

NBER WORKING PAPER SERIES

THE COST OF UNCERTAIN LIFE SPAN

Ryan D. Edwards

Working Paper 14093

<http://www.nber.org/papers/w14093>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

June 2008

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. This research was partially supported by NICHD grant T32 HD 07329. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2008 by Ryan D. Edwards. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Cost of Uncertain Life Span
Ryan D. Edwards
NBER Working Paper No. 14093
June 2008, Revised August 2010
JEL No. I10,J11,J17,O11

ABSTRACT

A considerable amount of uncertainty surrounds life expectancy, the average length of 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 less 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. But 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. Sustained reductions in the standard deviation of adult life span, which have largely ceased among advanced nations, accounted for about 15 percent of the total economic value of gains against mortality in the U.S. prior to 1950 but only about 5 percent since.

Ryan D. Edwards
Department of Economics
Queens College - CUNY
Powdermaker Hall 300-S
Flushing, NY 11367
and NBER
redwards@qc.cuny.edu

1 Introduction

Suppose that for a price, individuals could choose among several industrialized countries which would be their life-long home starting from birth. Infant mortality is negligible, and the only socioeconomic element that differs across regions is adult survivorship, which depends only on physical location.¹ If the choice is between living in country A, where life expectancy at birth, e_0 , is 80 years, or country B, where $e_0 = 78$ years, how much would an individual be willing to pay to be born and live in country A instead of B? Economists have developed an answer to this question based on how people respond to varying degrees of mortality risk. According to the review by Tolley, Kenkel and Fabian (1994), the average individual should value an extra year of life at about \$200,000 today after adjusting for inflation. Then we might expect that he or she might be willing to pay about \$40,000 today, or the present discounted value of those two extra year of life expectancy at birth, to live in country A rather than B.

In this paper, I argue that one needs to know more than just the average length of life to make this comparison. Suppose length of life is also more *uncertain* in country B, where there is a 20 percent chance of death before age 65, as opposed to only a 12 percent chance in country A. Put differently, suppose the standard deviation of length of adult life is 15 years in country B but only 13 in country B. How much would an individual pay to live in country A now? As I will show, lowering the standard deviation by two years should be worth about the same as attaining one additional year in mean length of life. So a reasonable answer is about \$60,000 in total to live in A rather than B, or an extra \$20,000 to avoid the heightened uncertainty in life span.

Valuing the spread in length of life to arrive at an answer like this could be done in three different ways. One could directly ask individuals about their preferences over volatility in life span. Alternatively, if there were a natural experiment that shifted the variance in length of life, one could examine revealed preferences. Finally, one can produce an estimate by calibrating a model and measuring the spread in a tractable way. Part of the problem with both of the first two options is that it is hard to conceive of a shock to length of life that shifts only the variance and

¹In reality, survivorship and socioeconomic status are closely interrelated. As I discuss in section 4.2, members of higher socioeconomic strata enjoy life spans that are longer on average and also less variable. But no amount of income or wealth can prolong life indefinitely or remove uncertainty altogether; all strata are subject to the cost of uncertain life span, albeit to varying degree.

not the mean.² Previous attempts at the first option have thus been limited in scope and results, and I know of no studies that follow the second. In this paper, I perform the last option, recovering a purely theoretical value for uncertainty in life span that is fully consistent with current economic thinking. It is similar in spirit to the ambitious theoretical work of Bommier (2006) but differs in that its focus is on empirical calibration. By revealing a large baseline estimate of the cost of uncertain life span, this paper sets the table for future studies of individuals' revealed preferences.

Choosing a convenient measure of the uncertainty is a key step. We can quantify uncertainty in human life span by interpreting the deaths in a life table as probability densities. Each panel in Figure 1 displays a column of the 1900 and 2000 U.S. life tables for both sexes combined presented by Bell and Miller (2005).³ The solid line in panel A displays the density function of life spans for the U.S. in 2000, which is a skew-left distribution around a mode of 85 years with a small spike at infant mortality. The unconditional mean of this distribution, also called period life expectancy at birth, or e_0 , is about 77 years.⁴ As is plainly visible in panel A, there is considerable variation around the mean even if we omit infant mortality, which is fixed at age 0. 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.⁵ It will also turn out to be conveniently tractable in the model I introduce in this paper. In 2000, S_{10} was about 15 years in the U.S.

Panel A in Figure 1 also reveals massive temporal change in this distribution. In 1900, life was much shorter on average, with $e_0 = 48$, and infant mortality was considerably higher. We can also

²The estimates of the value of a life year examined by Tolley, Kenkel and Fabian (1994) and others typically derive from economic responses to mortality shocks that change both the mean and the variance simultaneously, usually an on-the-job risk of death that is compensated with a market wage. If these shocks are small, they reveal the expected discounted value of a life year conditional on both the mean and variance inherent in the typical survivorship schedule.

³Data prior to 1933, the first year when all 48 states were included in the official Death Registration Area (Haines, 2001), are less representative but still suggestive of early patterns.

⁴The mean length of life conditional on survival to any early age past infancy, say age 10, is not much different at $M_{10} = e_{10} + 10 = 78$. This is because infant mortality is relatively low.

⁵As Edwards and Tuljapurkar (2005) explain, the standard deviation conditional on reaching practically any age greater than or equal to 1 but less than say 25 is a stable measure of adult inequality in advanced countries today. 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. When high, childhood mortality strongly affects S_1 and S_5 , while young adult mortality can be missed entirely by S_{25} . To be sure, opinions vary regarding the ideal type of measure of variability in 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. Preferable qualities of S_{10} and the IQR over the Gini are that they are relatively invariant to trends in infant mortality, which is etiologically distinct from adult mortality, and that they are invariant over additive translations of length of life. In a related paper, Edwards (2008) argues that the latter characteristic is desirable given the standard utility-theoretic model of preferences used in this paper. Trends in the IQR and S_{10} among industrialized countries over the last several centuries, as examined by Wilmoth and Horiuchi (1999) and Edwards and Tuljapurkar (2005), are essentially the same. As will become apparent in this paper, a considerable advantage of S_{10} is its analytical tractability.

see that adult life spans were much more uncertain than they are today, with $S_{10} = 24$, almost 10 years or some two-thirds higher. Figure 2 charts progress against uncertainty in adult life span, S_{10} , beneath gains in average life span, e_0 , over the last 150 years in Sweden, where historical statistics are of high quality.⁶ Wilmoth and Horiuchi (1999) and Edwards and Tuljapurkar (2005) investigate in greater detail the long-term trends in variability among advanced countries. Monumental declines in variance were achieved prior to the 1950s during the demographic and epidemiological transitions. Reductions have largely stalled since then, even as life expectancy continues to increase. Today, industrialized nations appear to be stuck with some uncertainty in life span, part of which is surely inherent to living creatures and thus unavoidable. But as Edwards and Tuljapurkar reveal, there are large and persistent differences today between high-variance countries, such as the U.S. and France with S_{10} around 15 years, and low-variance countries like Sweden and Japan, who enjoy levels of S_{10} near 13.

There are also large differences in S_{10} between rich and poor countries, although poor data quality among the latter group reduces the precision of inferences. Table 2, to be discussed later, shows that S_{10} has averaged around 20 years in Sub-Saharan Africa, where demographic and epidemiologic transitions have not completed and infectious disease remains a leading cause of death. In particular, the HIV/AIDS epidemic has increased variance in Africa beyond what it would normally be at comparable stages of development. Countries still in the middle of their demographic transitions but less hard hit by HIV/AIDS enjoy considerably lower levels of S_{10} , around 16.5 in the case of India and 14 in China, where population aging is commencing. While infant mortality is negligible in most advanced economies, it has been and remains very high in developing regions. As shown in Table 2, the chance of survivorship to age 10 in Sub-Saharan Africa is estimated at only 82.7 percent today, compared with 99.2 percent in high income countries.

Are there welfare costs associated with S_{10} and infant mortality? What are the implications of this broad definition of health inequality for assessing our progress against mortality, and for gauging the value of continued progress?⁷ Answering these questions requires a theoretical framework, and

⁶As discussed by Edwards and Tuljapurkar (2005), the trend toward lower variance is best characterized as a one-time, if extended, event during the half century of epidemiological transition in industrialized countries, while increases in life expectancy appear to be continuing apace. The decline of infectious disease as a leading cause of death during the early part of the 20th century not only raised life expectancy but lowered the variability in adult life spans considerably. After 1950, progress against chronic degenerative diseases like cancer and cardiovascular disease appears to have shifted the survivorship curve outward rather than compressing it (Wilmoth, 2003).

⁷To be sure, population S_{10} effectively conflates inequality as we would traditionally understand it, i.e., a difference

I explore the cost of life-span variance using a standard model of intertemporal choice. Both model and exercise are similar to those of Lucas (1987) in his classic assessment of the welfare cost of business cycles, but the nature of the uncertainty I consider here is quite different. A related paper (Edwards, 2009a) considers an isomorphic problem that is more directly analogous to Lucas's: the cost of cyclical fluctuations in mortality. Here, I am concerned with the cost of temporally static uncertainty in length of life. Economically large variation in S_{10} across groups also raises an entirely different and interesting set of research questions concerning behavioral aspects. Potential causes and consequences of uncertainty in length of life range broadly and are an interesting focus for future efforts. I discuss extant findings in this nascent area of inquiry but leave deeper questions to future research efforts.

As I show in this paper, the standard model of intertemporal optimization implies that individuals should be risk averse over life span, and the coefficient of absolute risk aversion in life span is approximately the time discount rate. Since average life span seems to be increasing linearly over time while the standard deviation around adult ages is now roughly fixed, constant absolute risk aversion is consistent with stable risk premia and thus is an intuitive result.⁸ As I discuss below, risk aversion over life span fits the relatively scant empirical evidence on stated preferences and behavior in medical settings. Bommier (2006) presents a more detailed and general modeling of risk aversion over life span that relaxes the assumption of additively separable preferences via another curvature parameter, which he does not calibrate. If risk aversion is indeed greater than implied by time discounting alone, as Bommier's framework assumes implies, then my estimates will understate the true cost of life-span uncertainty. To be sure, much research on consumption and asset pricing suggests that additive separability is too restrictive to explain many phenomena (Epstein

in mean life span between groups, with uncertainty in life span within groups or overall. While it is a measure of total *ex post* inequality, labeling it as such may be misleading. As Edwards and Tuljapurkar (2005) reveal, population-level S_{10} masks substantial differences in both mean and variance in life span across population subgroups, which indicates that inequality, uncertainty, and inequality in uncertainty are all at work.

⁸The intuition derives from an analogy to the relationship between time trends in aggregate consumption, risk premia in consumption, and preferences over consumption, a relationship 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.

and Zin, 1989, 1991; Browning, 1991). But Nordhaus (2003), Becker, Philipson and Soares (2005), Murphy and Topel (2006), Hall and Jones (2007), and others in the subfield on valuing longevity extension all assume additive separability. 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.

For reasonable parameter values, the standard intertemporal model suggests that variance in adult life span is quite costly, even when wealth can be fully and costlessly annuitized. Time discounting also makes infant mortality extremely costly, considerably more so than is implied by the sizable effect of infant mortality on mean length of life alone. There are several implications of my results. Each year in adult standard deviation is worth about half a year in the average; that is, an individual would agree to give up half a year in mean life span to obtain a standard deviation one year lower. In light of this, the value of differences in population health between low-average, high-variance countries and high-average, low-variance countries is larger than the differences in life expectancy alone. As I will show, the vast decreases in the standard deviation in the U.S. during the 20th century represent roughly 15 percent of the total value of gains against mortality. Previous estimates of the total economic value of mortality decline since 1900 (Nordhaus, 2003; Murphy and Topel, 2006) typically account for all changes in the survivorship schedule, whether involving the mean, variance, or higher moments, and thus are not biased downward. But my findings reveal that a significant amount of the gains was due to declines in variance, which have now largely ceased (Edwards and Tuljapurkar, 2005).

In valuation studies where only life expectancy at birth is considered due to data limitations, the omission of higher moments will bias results because variance is costly. In the final section, I revisit the efforts of Becker, Philipson and Soares (2005), who reveal much global convergence in world average “full income,” a measure of human well-being that incorporates the value of life along with income per capita. Due to lack of data, their analysis accounts for the average but not the variance in human life span. Incorporating new work by Edwards (2009*b*), I show that accounting for life-span uncertainty amplifies patterns of convergence in full income since 1970 for two reasons. Methodologically speaking, accounting for the presence of any variance, adult or otherwise, increases the measured value of a given improvement in life expectancy, and the

value is increasing in the variance. An observed gain in the mean length of life is thus worth more to high-mortality countries, which also tend to be poor. In addition, there has been much global convergence in infant mortality, more so even than in life expectancy (Moser, Shkolnikov and Leon, 2005), although considerably less in S_{10} . Infant mortality can be interpreted as the most costly type of variance in total life span, and reductions strongly increase full income even if life expectancy were held constant. Thus while higher mortality and variance in length of life still disproportionately penalize the poor, the value of improvements in mortality works to reduce global inequality in average well-being.

The structure of the paper is as follows. In Section 2, I contextualize this research by discussing extant work on choice under uncertain survivorship and literature on the value of life and on bequests. Section 3 presents the utility-theoretic model of the cost of uncertain length of adult life, compares analytic approximations of the central result to numerical simulations of the full model, and discusses the cost of infant mortality. In Section 4, I explore the implications of these insights, such as for decomposing the value of mortality declines and assessing convergence in average human well-being. The final section concludes.

2 Background

A wealth of research exploring the valuation of life has produced many insights regarding the willingness to pay for mortality reduction (Rosen, 1988; Viscusi, 1993; Johannsson, 2002; Aldy and Viscusi, 2003). Some investigators have examined the willingness to pay over historical periods and across geographical boundaries (Viscusi and Aldy, 2003; Costa and Kahn, 2004), while others use this evidence to value long-term increases in life span (Cutler and Richardson, 1997; Nordhaus, 2003; Becker, Philipson and Soares, 2005; Murphy and Topel, 2006). Most of these efforts model survivorship probabilities realistically, i.e., with implicit variance around mean life span. Panel B in Figure 1 displays the survivorship probabilities, or the cumulative of the density in panel A. But Becker, Philipson and Soares (2005), who examine a broad cross section of 96 rich and poor countries in 1960, 1990, and 2000, could not examine complete survivorship curves because the data were not available. Rather, they were forced to use only life expectancy at birth, the average of the distribution.

To be sure, higher moments do not matter if there is risk neutrality over length of life, and the degree of risk aversion is not trivially clear and could be zero. 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 decade 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.⁹ In a theoretical contribution, Bommier (2006) models preferences over life span with an extra parameter producing risk aversion even when the rate of time preference is zero.

But 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

⁹The 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 wide-ranging preferences they observe.

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 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 reveal the theoretical cost of uncertain life span in the standard intertemporal model common to economics, paying special attention 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, but I express the cost of variance in life span in terms of the value of the mean. This lessens the sensitivity of my results to the assumption of no intended bequests, because a bequest motive would work to reduce both simultaneously.

3 Modeling the cost of life-span uncertainty

Grossman (1972), Ehrlich and Chuma (1990), and others have modeled preferences over health capital and life extension by adapting standard economic frameworks of intertemporal choice. Here I continue that tradition but with a simpler model without health capital in order to enhance analytical tractability. It turns out that when future periods are discounted and additively separable, expected utility maximization 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.¹⁰ Lifetime expected utility is the sum of period utilities drawn from consumption, $u(c(t))$, weighted by the force of time discounting, $e^{-\delta t}$, and the probability that the individual is alive, $\ell(t)$:

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

The survivorship function, $\ell(t)$, shown in Panel B of Figure 1, is one minus the cumulative density function of life span. The decrement to $\ell(t)$ is the probability density function of life span, shown in Panel A and commonly called life-table deaths. Finally, we define the life-table probability of dying between t and $t + 1$ as $q(t) = \log[\ell(t)] - \log[\ell(t + 1)]$.¹¹ Panel C of Figure 1 shows how mortality tends to increase exponentially with age, as originally revealed by Gompertz (1825). For simplicity, I do not explicitly model health status. To the extent that younger, healthier life-years are probably more valuable than older, sicker life-year years, my model will tend to underestimate the true cost of a mean-preserving spread in life span.

¹⁰In reality, where mortality prospects vary across types, we might posit that an individual begins life under a “veil of ignorance,” facing the full level of population uncertainty, but then learns his or her type later, which lowers uncertainty. This would alter behavior once the type became known. I do not model this dynamic here for two reasons. Parents, who possess better forecasts, make most of the decisions for children. Edwards and Tuljapurkar (2005) show that although there are interesting differences in life-span uncertainty across U.S. subgroups, they are not enormous.

¹¹This form is more convenient for expositional purposes here, but life-table $q(t)$ is typically defined implicitly as $\ell(t + 1) = \ell(t)[1 - q(t)]$, in order to attrit the entire cohort at a finite age. In continuous time, $q(t)$ is the hazard or mortality rate.

Suppose the individual has a financial endowment W that can be consumed or saved at a fixed market rate of interest, r . For simplicity, there is no labor, education, capital, or financial risk in this model.¹² In a market without annuities, the budget constraint requires the individual to finance the present value of all future consumption out of wealth:

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

Under this budget constraint, the model will produce unintended bequests whose size varies inversely with length of life. This is because the individual must engage in a type of precautionary saving against the risk of living too long, and unused savings become bequests. If instead actuarially fair annuities are available, the budget constraint takes on a different form:

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

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.¹³

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

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

where D is an indicator of the lack of annuities. Under full annuitization, $D = 0$ and mortality cancels out of equation (4) because survivorship appears in both the objective and the constraint.

¹²Kalemlı-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.

¹³In 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.

Other things equal, consumption will tend to 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). This will produce a consumption trajectory that looks like the survivorship curve.

We are interested in the expected utility cost of variance in life span, or equivalently of facing a mortality schedule $q(t)$ that is rising exponentially rather than staying at zero until rising to 1 at the date of death. Based on equation (4), it is tempting to assert that annuities must offset the cost of life-span variance, since $q(t)$ does not appear in the Euler equation under full annuitization. In fact, annuitization reduces the cost of variance but cannot eliminate it, as I will now show.

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 use numerical simulations to relax assumptions about annuities and the distribution of life spans. Since period utility is a function of consumption, which typically depends on time or age and thus the distribution of life span, it is convenient to assume a specific functional form of utility in order to proceed. Power utility is a standard assumption in economics and a reasonable baseline. Since my qualitative results are driven by discounting and time-separability, they are likely to generalize to other period utility functions. More fundamental questions, such as whether preferences actually reflect time-separability and exponential discounting, and whether they are age independent, are perfectly valid but beyond the scope of this paper to address.

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 .¹⁴

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

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

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

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.¹⁵

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

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

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 , we could

¹⁴To 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.

¹⁵Mean 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 lifetime 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 satisfying the budget constraint. Numerical simulations reveal that the effects on $c(0)$ of changing M or S^2 are small.

use equation (6) to solve the integral and then find $c(0)$ by taking the expectation.¹⁶ But we are 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):

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

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

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

where K is the constant utility of being alive, and

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

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

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 . 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.¹⁷

We would expect individuals who are risk averse over life span to be hurt by uncertainty in life span, and this is exactly what we see if $T \sim N(M, S^2)$.¹⁸ When life spans are normally distributed,

¹⁶Without annuities, the presence of $q(t)$ in the consumption function precludes analytical solutions because mortality increases exponentially with age.

¹⁷As discussed earlier, constant absolute risk aversion over life span seems consistent with roughly linear time trends in life expectancy and a stable S_{10} during the era since World War II. This is because constant relative risk aversion over consumption risk, such as implied by power utility, fits the pattern of exponential growth in consumption and stability in the standard deviations of log consumption and of asset returns, and in financial risk premia (Campbell and Viceira, 2002).

¹⁸As 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

expected lifetime utility in this model is

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

by virtue of the properties of lognormality. That is, 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.¹⁹

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:

$$p_S = \frac{\partial EU / \partial S}{\partial EU / \partial M}. \quad (12)$$

When the utility shifter K is nonzero, this ratio is mathematically complicated.²⁰ A first-order Taylor expansion around $K = 0$ reveals

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

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

$$p_S = -\hat{\delta}S. \quad (14)$$

The price of a standard deviation in life span, p_S , is negative when $\hat{\delta} > 0$ because S is a bad. That is, 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

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.

¹⁹It 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.

²⁰The 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 $\hat{\delta}$.

$\hat{\delta}$, 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 willingness to bear additional risk falls with the level of risk because its marginal disutility rises. 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.²¹

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

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

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.

We could also solve for the dollar price of S by constructing the ratio of marginal lifetime utilities of S and $c(0)$ or wealth instead, but pricing the variance in terms of mean life span is useful for two reasons. If bequests are intended, they should attenuate both the marginal disutility of life-span variance and the marginal utility of mean life span, while probably amplifying the marginal utility of consumption, since your money could buy happiness through your children's happiness. Although evidence suggests that bequest motives must not be very strong, the price of S in terms of M is likely to be a more stable relationship than the price of either in terms of money. The latter depends on assumptions about the curvature of utility in terms of consumption, given by γ here, and thus the elasticity of intertemporal substitution in consumption, and if they are not separated, the coefficient of relative risk aversion in consumption. As I will discuss later,

²¹According 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.

conceptualizing the price of variance in terms of the mean rather than consumption is also useful for interpreting current cross-national and intertemporal differences in population health.

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 $S = 15$, 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

I next examine the sensitivity of the analytical result, $p_S = -\hat{\delta}S$, to alternative assumptions about wealth annuitization and mortality using numerical methods. 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. For better tractability and for clearer comparisons, I begin by modeling life span as normally distributed.²² For now I omit infant and child mortality and set S equal to the standard deviation in life spans above age 10, S_{10} . Similarly, for M I use M_{10} , which is also equal to remaining life expectancy at age 10, e_{10} , plus 10. I set mean and variance equal to U.S. levels in 1994, $M_{10} = 76.85$ and $S_{10} = 15.66$ years, and I search for the mean life span that compensates expected utility for a decrease in S_{10} to 15.05, the 1999 level. 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.

²²I 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 distributed with means around age 80, densities past age 110 are not miniscule. Truncating at age 110 actually creates significant skewness in the distribution, and probably changes the mean and variance. Such skewed distributions actually produce a p_S locus that better resembles that under realistic survivorship because real life spans are skew-left.

3.3.1 Normally distributed life spans

Figure 3 plots 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. I fix the rate of market interest at $r = 0.03$ and examine how varying δ , shown along the horizontal axis, changes p_S given $S_{10} = 15.66$. The steeper, concave line shows the locus obtained from the numerical model with full annuitization of wealth, while the lower dashed line depicts the schedule that results from the numerical model without annuitization.

The two solid lines reveal relatively limited differences between the analytical and numerical versions of the model with annuities. They cross at $\delta = 0.03$, with the numerical model producing a more steeply sloped p_S locus that becomes concave at high δ . These differences are due to the presence of the utility shifter K , which when negative tends to increase p_S when δ is large, as shown by equation (13). When $K > 0$, the numerical locus is flatter than the analytical locus, and when $K = 0$, the lines overlap. Both lines show that when $0 < \gamma < 1$ and δ is small relative to $r = 0.03$, $\hat{\delta}$ can actually become negative, which shifts p_S positive. 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 relationship in Figure 3 between the cost of uncertainty under annuitization and the cost without annuitization is more interesting. The dashed line, which shows p_S without annuities, is significantly lower than the other two, indicating that life-span variance is more costly at any δ without annuities. At the baseline of $r = \delta = 0.03$, p_S without annuitization is about -0.75 , more than half again as large as p_S with annuitization, which is about -0.45 . This is an intuitive result insofar as annuitization is designed to remove risks to consumption associated with uncertain life span. That annuitization removes only a little over one third of the total cost of life-span uncertainty under baseline parameter values is more surprising. In this model, the direct utility cost of S is more important than the cost of consumption uncertainty.

One can also vary r while holding $\delta = 0.03$ fixed, which is shown in Figure 4. Here, an even greater difference emerges between the numerical and analytical models under annuitization, which again turns out to be linked to the utility shifter K . The numerical model reveals considerably

more negative p_S when $r < \delta$ and $K < 0$. This is because when $r < \delta$, consumption optimally declines over time. If in addition $K < 0$, the benefit of living long is weakened considerably relative to the cost of dying early, so the cost of S is higher. As before, we see that p_S is more negative when annuities are not available, regardless of the level of r . Each locus also shows that the cost of S is declining in magnitude with r . If $0 < \gamma < 1$, then a very large r , around 0.15 with these parameters, can produce a negative $\hat{\delta}$ and thus a positive p_S (not shown). As before, when r is high relative to δ , the price of future consumption is low. Intertemporal substitutability is high when $0 < \gamma < 1$, and the individual will substitute toward future consumption. If r is high enough, a mean-preserving spread in life span might actually improve expected well-being for someone so dependent on future utility.

The period utility curvature parameter, γ , does play a role here but is much less interesting than one might expect.²³ We are accustomed to thinking of γ as the coefficient of relative risk aversion over gambles in consumption, since a higher γ represents a more concave period utility function. Here, its influence is more easily understood as deriving from its other role, the inverse of the elasticity of intertemporal substitution.²⁴ 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 $\hat{\delta} \approx \delta$.

²³As I discussed in note 14, 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 important. Functionally, the MRS between average length of life and the variance depends mostly on the time discount rate. Intuitively, it is only utility curvature over length of life, not consumption, that matters in the cost of life-span variance. The relevant issue here is whether additively time-separable preferences are overly confining, and I defer to established tradition in this subfield while acknowledging it remains an open and interesting question.

²⁴An interesting parallel emerges but is left for future research. We have seen how δ , the rate of time discounting, is also the coefficient of absolute risk aversion over life span in this model. It is an open question whether attitudes toward time and risk as regards periods of life ought to be or even can be separated, as they have been with attitudes regarding consumption (Epstein and Zin, 1989, 1991).

3.3.2 Realistic adult life spans without infant mortality

Modern distributions of human life span are skew-left and leptokurtic, not normal. The skewness is of direct interest here, because it implies that a mean-preserving spread in life span lowers survivorship probabilities asymmetrically. When there is leftward skewness, an increase in variance reduces survivorship at young adult as well as adult ages, or across two age groups, while increasing it only at old ages. Because of discounting, skewness should amplify the cost of uncertain life spans.

It is tricky to model realistic life spans with particular means and variances because we do not have a convenient functional form of the 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 $M_{10} = 76.85$ and $S_{10} = 15.66$.²⁵

Figure 5 is an analogue of Figure 3 with realistic adult mortality. 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 as before, since I cannot mathematically model realistic life spans. The other two schedules reveal levels of p_S that are more negative than in Figure 3, especially for larger δ . That is, the cost of life-span uncertainty is indeed higher when I model life span realistically, with leftward skewness that places a wider range of early years at risk. The thin solid line, representing the output of the numerical model with annuities, shows $p_S \approx -0.7$ when $\delta = 0.03$, which is considerably greater in magnitude than the $p_S = -0.45$ from the simple model. As δ rises, both numerical models show more precipitous increases in the cost of variance when adult life span is realistic.

²⁵I 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.

3.3.3 Fully realistic life spans with infant mortality

Some human life spans end at birth or shortly thereafter, and the variance attributable to infant and child mortality is by far the most costly. One less infant death is one more average adult life span, other things equal, so the value of a marginal reduction in deaths at age 0 should be equal to the reduction times the expected utility of an entire average life span. As we have seen, the marginal cost of adult variance is much lower, worth a fraction of an additional year of life rather than an entire life span. Infant mortality is relatively uncommon in advanced countries today, so ignoring it when focusing on the differences between Sweden and the U.S. in 1999 is perfectly acceptable, as I will demonstrate. But as will become clear in Section 4, infant mortality plays a very important role in understanding the value of mortality declines in rich countries over longer periods of time and in poor countries, where infant mortality remains high.

Here, I am concerned with how the presence of infant mortality may affect p_S . Given a fixed pattern of adult survivorship, higher infant mortality reduces the average length of life and raises total variability. By making mean life scarcer and raising the marginal utility of M , the presence of any infant mortality should lower p_S . But because the marginal disutility of adult variance rises with adult variance, infant mortality could conceivably amplify the marginal disutility of S if the latter were also sensitive to total life-span variance. This would tend to increase p_S . To ascertain which of these countervailing effects wins out, 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 that offsets the reduction in S from 15.66 to 15.05.²⁶

Figure 6 depicts p_S as functions of δ when $r = 0.03$. Survivorship is now fully realistic, with a fixed amount of infant mortality in addition to leftward skewness and leptokurtosis in the adult hump. Including infant mortality does not change the qualitative results, but it does attenuate p_S compared to Figure 5, where adult mortality is realistic but there is no infant mortality. The presence of infant mortality apparently increases the marginal utility of mean life more than it raises the marginal disutility of a standard deviation in life span, if it changes the latter at all. The magnitude of p_S may be less with fully realistic mortality than without infant mortality, but the

²⁶I 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 cdf sums to unity.

analytical solution remains a conservative estimate of the true cost.²⁷

3.4 The cost of infant mortality

A separate but related question 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 longer periods of time or wider geographic space, so it could easily prove important for assessing well-being more broadly.

Unlike a central moment like S_{10} , infant mortality strongly affects the mean length of life, especially when infant mortality is high. It is partly for this reason why model life tables and estimates of period life expectancy at birth in developing countries are primarily based on measures of infant mortality and often little else. Given this, one might expect differences in period life expectancy to sufficiently encompass differences in infant mortality so that valuation studies need only consider the former. But the question is how each affects lifetime utility, and in fact they can operate very differently. Here, I treat infant utility as if it were simply shortened adult utility. This is consistent with its treatment in the literature on the value of life extension, but it may not be a valid assumption in a more realistic setting such as with parental altruism, evolving time preference, and so on.

Assume for simplicity that infant mortality is completely described by $q(0)$, the probability of death between ages zero and one, and suppose $q(0)$ rises by some amount $i \in (0, 1)$. Under these circumstances, a decline in infant mortality of $i = -0.01$ raises survivorship, life expectancy at birth e_0 , and lifetime expected utility each by 1 percent.²⁸ By comparison, proportional increases in e_0 brought about solely by reductions in adult mortality typically produce less than proportional

²⁷To 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.

²⁸When 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 $\ell(t)$.

increases in lifetime utility.²⁹ 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 implied by their effect on life expectancy at birth alone. This insight will prove useful in Section 4, when I reassess global convergence in mortality after accounting for adult variance and infant mortality in addition to mean length of life.

For another perspective on the cost of infant mortality, I can examine p_i , the cost of infant mortality in terms of the mean, the analogue of p_S . Assuming complete annuities, normally distributed life spans, and $K = 0$ for simplicity,

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

where I have canceled out period utility. As with p_S , the price of a bad is negative. Because the shock i multiplies survivorship and is effectively the change in the probability of living at all, the marginal lifetime utility of a unit change in i is very large: all of lifetime expected utility. This will always exceed the marginal utility of a mean life year, which is the discounted value of a single year lived at the mean length of life. Equation (16) states that a unit reduction in i is preferable to a unit increase in M , which is 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. Whenever M is sufficiently greater than 1, of course it is more valuable to live at all than to increase the mean length of life.

A more interesting measure is $p_i/100$, the price of one percentage point in the infant survival probability in terms of the mean length of life. Recasting $\hat{\delta}$ in terms of percentage points in equation (16) reveals that when $\delta = 0.03$, p_i is about one third of the risk-adjusted return to living M years. This reaches and begins to exceed 1 in absolute value around $M = 40$ when there is realistic adult variance. That is, when life expectancy exceeds 40, a decline in infant mortality of a percentage point is more valuable than an additional year in the mean. Life expectancy at birth exceeds 40 in the majority of countries during the modern period, but notable exceptions include countries

²⁹For 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.

hard-hit by HIV/AIDS.

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. The main issue is measuring it, since it is a latent preference parameter. 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). The baseline of $\delta = 0.03$ therefore certainly seems reasonable by financial market standards. 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.

But discount rates appear to vary over individuals within and across countries (Barsky et al., 1997; Becker and Mulligan, 1997), as do life spans. We must be careful about assessing the costs of different life-span variance under different discount rates if time preferences and life spans are related to one another. Fuchs (1982) views the discount rate as determining health investments, while Becker and Mulligan (1997) see wealth, uncertainty, and health or the length of life as jointly determining the discount rate.

If the causality runs from life span to the discount rate, then a reasonable way to deal with the endogeneity is by treating the discount rate as a universal, unvarying parameter set at some reasonable baseline and used to price variance, which may differ. If causality runs the other way, the story becomes more complicated. My model shows that a higher discount rate increases the costliness of life-span variance, which should incentivize behavior leading to lower variance. While the net effect on p_S is ambiguous, the true cost of variance will be higher than measured with a standard discount rate. But it is also plausible that higher discount rates could produce more life-span variance. This could occur if myopic, risky behavior carries long-term costs but yields some short-term benefits, which I do not measure. If this were the case, we might find that the

costs of additional life-span uncertainty were outweighed by the benefits, clearly a problematic but also a seemingly far-fetched scenario. Public health campaigns certainly suggest that most risky behaviors are socially undesirable, if they ultimately say little about net private returns. A normative assessment remains elusive here, a point I concede but relegate to future efforts.

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 a difference in p_S of almost 0.1 year if $\delta = 0.03$, an increase in the costliness of life-span uncertainty of 20 percent.³⁰

With this much subgroup variation, one must be 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. An increase in life expectancy among a high-mean, low-variance group would technically increase aggregate S_{10} but also be Pareto-improving, provided individuals knew their group. 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 (Meara, Richards and Cutler, 2008) suggests between-group inequality is indeed increasing. The implications are unclear, especially because within-group inequality should be decreasing if aggregate S_{10} is indeed stable. 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; period S_{10} thus measures uncertainty for a fictitious cohort of individuals living their entire lives in a single period. Although period S_{10} is the appropriate variance analogue of period life expectancy, the most commonly cited

³⁰The less fortunate also bear a heavier burden if they have disproportionately less access to annuities, which we have found to offset perhaps one third of the cost of life-span uncertainty. This characteristic may be observationally linked to high δ , since one reason why low-SES individuals might appear to have high δ , myopia, or insufficient saving, is if they face liquidity constraints or incomplete markets. We would expect that access to annuities markets are also poor for liquidity constrained individuals.

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 7. I plot both measures for the U.S. since 1900 using annual period life tables from the Human Mortality Database (2009) and cohort life tables collected and forecast by the Social Security Administration (Bell and Miller, 2005). 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 face $S_{10} = 15$, face drastically less uncertainty in their life spans than those born around 1900, for whom $S_{10} = 21$.³¹

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. More generally, the chief insight of this paper, that variance in length of life is costly under time discounting, leads us to reassess what we know about the value of mortality improvements over time and across countries. 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 $S_{10} = 15$ 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

³¹Technically, 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, 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 $S_{10} = 15.33$, or only 0.03 year higher than that implied by the median forecast life table.

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.³² 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.

To be sure, the U.S. and Sweden represent two relative extremes where differences between e_0 and S_{10} are very large. A slightly different perspective 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 certainly increases when we account for different variances, but the effect is moderate in size.

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 the U.S. over historical periods using estimates of the willingness to pay for mortality reductions. 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?

³²For 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 $M_1 - M_2 = \hat{\delta}(S_1^2 - S_2^2)/2$.

Murphy and Topel highlight the wide dispersion of gains across ages prior to 1950, when reductions in life-span variance were sustained and thus much more important. This pattern is evident in the top two rows of Table 1, which list life expectancy at birth (e_0) and the standard deviation in life spans above age 10 (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. Similarly, S_{10} fell by 8 years, from 24 to 16, a factor of one third. While life expectancy continued to increase steadily after 1950, rising an additional 8.3 years by 2000, at a rate of 0.166 per annum, further reductions in S_{10} were slight, totaling only 1.1 years 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 measure of the total gains against mortality in terms of mean life years. By this accounting, 14 percent of the total gains over the century were due to reductions in life-span variance, but after 1950, that figure falls to 5.6 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 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. (2001) 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 (2009*b*) 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.³³ As discussed by Edwards, these new data are of varying and often dubious 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

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

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

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

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:

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

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).

³³In 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.

As discussed in Section 3, high infant mortality also tends to reduce $A(\ell)$, holding other things including e_0 equal. Indeed, infant mortality can be conceptualized as additional variance in human life span that is highly non-Gaussian and non-central. 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 $e_0/(1-d) > e_0$, 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.³⁴ 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 (2009*b*) and income data from the Penn World Table, Maddison (2003), and the IMF. All regions gained income per capita during the interval,

³⁴In 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 implications for political stability and the continued spread of the demographic transition.

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 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, S_{10}^2 . 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 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 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 limited convergence in S_{10} , are surely also at work.

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, directly 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.³⁵ 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. A possible outlier based on this perspective, HIV/AIDS is something of a hybrid disease afflicting old and young alike.

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 variance should precede sustained progress in the adult mean or mode.³⁶

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.

³⁵A 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.

³⁶To 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.

Kalemli-Ozcan and Weil (2002) 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 seem to support this view. In contrast, Hurd, Smith and Zissimopoulos (2004) find very 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 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. While not a new subfield, the study of interactions between health, mortality, and economic decision making is likely to produce results that are of great interest to researchers and policymakers in many fields, in particular those of development and health inequality where mortality conditions vary the most.

5 Conclusion

In a standard model of time-separable utility with no bequest motive, 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. Even when wealth is fully annuitized, hypothetical individuals with these preferences are hurt by uncertainty in life span and would be willing to trade away $p_S = -\delta S$ years of mean life span in return for one less year in standard deviation.

If $\delta = r = 0.03$, the average American would be willing to give up 0.45 life year in return for one year less in the standard deviation at ages over 10, which is currently about $S_{10} = 15$. 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. In the thought experiment I posed in the introduction, country A enjoyed $M_{10} = 80$ and $S_{10} = 13$, while country B experienced $M_{10} = 78$ and $S_{10} = 15$, roughly like Sweden and the U.S. today. If the 2 year difference in average life span were worth \$40,000, my results suggest the 2 year difference in S_{10} should be worth about one additional life year, or another \$20,000. An unborn baby engaged

in birthplace arbitrage would be willing to pay \$60,000 more to live in Sweden. To be sure, my approximation of the price of uncertainty, $p_S = -0.45$, is just that, an approximation. Results vary with different assumptions regarding the relative levels of the real interest rate and the rate of time discounting, and with the characteristics of mortality. But as a rule of thumb, it is a useful and parsimonious representation of the key idea.

Babies choosing birthplaces would also be remiss not to consider infant mortality, to put it mildly. Although infant mortality is low in advanced countries today, 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 extremely valuable years. 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.

This framework allows me to assess the contribution of reductions in adult life-span variance and infant deaths to the overall value of progress against mortality. This decomposition is interesting because gains against adult variance and infant mortality have largely plateaued in developed countries, while increases in the mean continue apace. In the U.S., reductions in S_{10} prior to 1950 were relatively large, falling from 24 in 1900 to 16 by 1950, and they accounted for more than 15 percent of the total value of mortality declines. Since 1950, S_{10} has declined by only more year, contributing only about 5 percent to the total value of all progress, which is dominated by increases in the mean, albeit much decelerated themselves.

This paper also implies that health inequality is typically larger between countries than is implied by life expectancy alone, because the level of life-span uncertainty is often different as well. In particular, poor countries typically experience much higher levels of adult variance in length of life than do rich countries, in addition to their lower life expectancies. If we account for variance, global health inequality must be higher than we would perceive it based on the inequality in life expectancy alone. But in a dynamic setting, the importance of variance in length of life manifests itself somewhat differently. Purely methodological insights are that the presence of some variance increases the measured value of a given improvement in life expectancy, the value is increasing in the variance, and poor countries experience higher variance and thus larger gains from a given

change in life expectancy. Worldwide, moderate declines in adult variance combined with broad-based reductions in infant mortality, the larger source of unconditional variance in human life span, have produced much global convergence in average human well-being. These findings amplify those in the pioneering work of Becker, Philipson and Soares (2005), who due to data limitations could only consider life expectancy.

Several limitations constrain the scope of this paper. 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 the cost of 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 rather 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. I leave issues of the emotional costs of uncertain life span to future research in this nascent subfield. Many questions remain about time discounting and its relationship to risk preferences over periods of life, and I have made no attempt to investigate these issues here. Bommier (2006) posits a model of preferences over consumption and life span with a free parameter governing risk aversion over life span. 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. I intend my estimate of the cost of uncertain life span to be a provocative motivation for further research into this topic.

The 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. Since 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 life when uncertainty is high. A 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), a refocusing would not come without cost. A somewhat 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.
- 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.
- 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.
- Browning, Martin. 1991. "A Simple Nonadditive Preference Structure for Models of Household Behavior Over Time." *Journal of Political Economy* 99(3):607–637.
- 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. and Matthew E. Kahn. 2004. "Changes in the Value of Life, 1940–1980." *Journal of Risk and Uncertainty* 29(2):159–180.
- Cutler, David M. and Elizabeth Richardson. 1997. "Measuring the Health of the U.S. Population." *Brookings Papers on Economic Activity, Microeconomics* pp. 217–271.
- 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. 2008. "Measuring Socioeconomic Differentials in Mortality." Unpublished manuscript.
- Edwards, Ryan D. 2009a. "The Cost of Cyclical Mortality." *B.E. Journal of Macroeconomics*, 9(1):Contributions, Article 7.
- Edwards, Ryan D. 2009b. "Trends in World Inequality in Life Span Since 1970." Unpublished manuscript, May.
- 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.
- Ehrlich, Isaac and Hiroyuki Chuma. 1990. "A Model of the Demand for Longevity and the Value of Life Extension." *Journal of Political Economy* 98(4):761–782.
- 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.
- Gompertz, Benjamin. 1825. "On the nature of the function expressive of the law of human mortality and on a new mode of determining life contingencies." *Philosophical Transactions of the Royal Society of London* 115:513–585.
- Grossman, Michael. 1972. "On the Concept of Health Capital and the Demand for Health." *Journal of Political Economy* 80(2):223–255.
- Haines, Michael R. 2001. "The Urban Mortality Transition in the United States, 1800-1940." *Annales de Démographie Historique* 1:33–64.
- 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.
- Judge, Ken, Jo-ann Mulligan and Michaela Benzval. 1998. "Income Inequality and Population Health." *Social Science & Medicine* 46(4–5):567–579.
- Kalemli-Ozcan, Sebnem and David N. Weil. 2002. "Mortality Change, the Uncertainty Effect, and Retirement." *NBER Working Paper* 8742.
- 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.

- 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.
- 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., Joshua Salomon, Omar Ahmed, Christopher J. L. Murray and Doris Mafat. 2001. "Life Tables for 191 Countries for 2000: Data, Methods and Results." *GPE Discussion Paper Series* 9. Geneva: World Health Organization.
- Lucas, Robert E. 1987. *Models of Business Cycles*. New York: Blackwell.
- Maas, A. and P. 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 Centre.
- McNeil, B. J., R. Weichselbaum and S. 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 HYEes — 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.
- 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.

- 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.
- 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.

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}	24.0	16.0	14.9
Survivorship to age 10, ℓ_{10}	0.782	0.963	0.991
	1900-2000	1900-1950	1950-2000
[1] Average S_{10}	19.5	20.0	15.4
[2] Average p_S	0.58	0.60	0.46
[3] Change in S_{10} , ΔS_{10}	9.1	8.0	1.1
[4] Life year benefit of ΔS_{10} , [2] \times [3]	5.3	4.8	0.5
[5] Average ℓ_{10}	0.886	0.872	0.977
[6] Life year benefit weighted by average ℓ_{10} , [4] \times [5]	4.7	4.2	0.5
[7] Change in life expectancy, Δe_0	29.0	20.7	8.3
[8] Total improvement in life years, [6] + [7]	33.7	24.9	8.8
[9] Share due to ΔS_{10} , [6] \div [8]	0.140	0.169	0.056

Notes: 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 average number of years lived. S_{10} is the standard deviation of length of life conditional on survivorship to age 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.

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}	Survivorship ℓ_{10}	GDP per capita	Life exp. e_0	Life std. dev. S_{10}	Survivorship 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 is comprised of 180 countries including Taiwan, which is placed in the high income group. 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 (2009b). The value of life expectancy gains and full income are based on the author's calculations using the methodology and parameter values of Becker, Philipson and Soares (2005) and 1970 as the base year, and they are calculated three ways. 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.

Figure 1: The distribution of life spans, survivorship, and log mortality in the U.S. in 1900 and 2000




Source: Bell and Miller (2005) and author's calculations. These data are unweighted averages of sex-specific life-table entries. In panel A, the density of deaths at age 0 was 0.1328.

Figure 2: Average life span, e_0 , and the standard deviation of adult life span, S_{10} in Sweden since 1750




Source: Author's calculations and Human Mortality Database (2009).

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




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

Figure 4: The price of S , a standard deviation in life span, in terms of mean life span as a function of r when $\delta = 0.03$ and life span is normally distributed




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

Figure 5: The price of S , a standard deviation in life span, in terms of mean life span as a function of δ when $r = 0.03$, adult survivorship is realistic, but there is no infant mortality




Notes: The thick black line shows p_S , the price of a standard deviation in life span in terms of mean life span, as a function of r for the simple analytical model when the standard deviation of life span, $S = 15.66$, the level prevailing in the U.S. in 1994, and when $\delta = 0.03$. In the simple analytical model, when consumption is fixed, $p_S = -\delta S$. The thin solid line shows the locus in the full numerical model with complete annuitization of wealth. The dashed line depicts the locus in the model with no annuities. In this simulation, the survivorship weights are taken from a modified period life table for the U.S. in 1999, provided by the Human Mortality Database (2009). Life-table deaths at ages under 10 are set to equal deaths at age 10, and the entire distribution is rescaled to sum to unity.

Figure 6: The price of S , a standard deviation in life span, in terms of mean life span as a function of δ when $r = 0.03$ and survivorship is fully realistic, with infant mortality



Notes: The thick black line shows p_S , the price of a standard deviation in life span in terms of mean life span, as a function of r for the simple analytical model when the standard deviation of life span, $S = 15.66$, the level prevailing in the U.S. in 1994, and when $\delta = 0.03$. In the simple analytical model, when consumption is fixed, $p_S = -\delta S$. The thin solid line shows the locus in the full numerical model with complete annuitization of wealth. The dashed line depicts the locus in the model with no annuities. In this simulation, the survivorship weights are taken from a modified period life table for the U.S. in 1999, provided by the Human Mortality Database (2009). Life-table deaths at ages under 10 are set to equal deaths at age 10, and the entire distribution is rescaled to sum to unity.

Figure 7: The standard deviation in life span above age 10 in the U.S. by year and by birth cohort



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