

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 David Canning, Shripad Tuljapurkar, Neil Bennett, Ronald Lee, Deborah Balk, 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.

© 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
JEL No. I10,J11,J17,O11

ABSTRACT

A considerable amount of uncertainty surrounds life expectancy, e_0 , the average length of life. The standard deviation in adult life spans, S_{10} , is about 15 years in the U.S., and theory and evidence suggest it is costly. In this paper, I calibrate a standard intertemporal model to show that one less year in standard deviation is worth about half a mean life year. Differences in S_{10} amplify measured differences in e_0 between the U.S. and other industrialized countries, and accounting for historical gains against S_{10} raises the total value of mortality declines during the last century by about 25 percent.

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, you could choose which among several industrialized countries would be your 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 where you live.¹ If you had to choose 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 you 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. The average individual, who values an extra year of life at about \$200,000, would be willing to pay about \$40,000 today, which is the present discounted value of those two extra year of life.

In this paper, I argue that you need to know more than just e_0 to get your money's worth. Suppose I told you that your life span is also more *uncertain* in country B, where there is a 20 percent chance you will die before age 65, as opposed to a 12 percent chance in country A. Now how much would you pay to live in country A? As I will show, a reasonable answer is about \$60,000 in total, 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 can be done in two different ways. Either I could directly ask individuals about their preferences over volatility in life span, or I could produce an estimate through calibrating a model, measuring the spread in a tractable way. This paper discusses early attempts at the former, which have been limited in scope and results, and performs the latter, arriving at a purely theoretical value for uncertainty in life span that is fully consistent with current economic thinking. It is similar in spirit to the 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

¹In reality, survivorship and socioeconomic status are closely interrelated, of course. 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.

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).² 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.³ 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. 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 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.⁴ Today, industrialized nations appear to be stuck with some uncertainty in life span, part of which in any event is undoubtedly unavoidable and inherent to living creatures. But as Edwards and Tuljapurkar

²Data 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.

³The 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.

⁴As 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).

(2005) reveal, there are significant differences between high-variance countries like the U.S. and France and low-variance countries like Sweden and Japan. Furthermore, developing countries that have not yet completed their demographic transitions surely still suffer from higher variance in life span, although good data are scarce. Is there a welfare cost associated with S_{10} ? 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?⁵ Answering these questions, which I leave to future efforts, requires a theoretical framework.

In this paper, I explore the cost of life-span variance using a model of intertemporal choice that is standard in economics. 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. In a companion paper (Edwards, 2007), I consider a closely related topic 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.

I find that 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 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. As I discuss below, risk aversion over life span also fits the available empirical evidence on stated preferences and behavior in medical settings. Bommier (2006) discusses more general modeling of risk aversion over life span that is independent of the discount rate, but he does not calibrate the extra parameter. If risk aversion is indeed greater than implied by time discounting, my estimates understate the true cost of life-span uncertainty.

⁵To be sure, population S_{10} effectively conflates inequality as we would traditionally understand it, i.e., a difference 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 there is inequality, uncertainty, and inequality in uncertainty all at work.

For reasonable parameter values, the model suggests that variance in adult life span is quite costly even when wealth can be fully and costlessly annuitized. Each year in standard deviation is worth about half a year in the average; that is, an individual would agree to give up half a year in average life span to obtain a standard deviation one year lower. Accounting for the vast decreases in the standard deviation in the U.S. during the 20th century adds roughly another 25 percent to the total value of mortality reductions of the type estimated by Nordhaus (2003) and Murphy and Topel (2006), who weight age-specific gains with survivorship probabilities, implicitly assuming risk neutrality and thus no benefit from declining variance. Another implication is that the high levels of life-span uncertainty persisting in developing countries today is likely to temper the optimism of Becker, Philipson and Soares (2005), who find much global convergence in a measure of “full income” that accounts for the average but not the variance in life span.

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).

Many 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 density function of life span corresponding to the pdf in panel A. But Becker, Philipson and Soares (2005) cannot, since the data do not exist. They are forced to use life expectancy alone. Even when survivorship schedules are available, measuring total

progress as the survivorship-weighted sum of the value of gains in life years is equivalent to assuming risk neutrality over life span.

To be sure, the degree of risk aversion over years of life is not trivially clear, and it 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.⁶ 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. Annu-

⁶The 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.

ities and life insurance are two examples of market instruments that diversify risks associated with uncertain life spans, while bequests are a nonmarket instrument that could also conceivably hedge against life-span uncertainty. In this paper, I show that while full annuitization removes all *consumption risk* associated with uncertain life span, and therefore improves welfare, annuities cannot remove the *utility risk*. Even under full annuitization, life-span uncertainty is costly. Whether life insurance helps offset S_{10} is more difficult to say, since it affects individual utility only through the bequest motive. That is, actuarially fair life insurance is like a precommitted bequest, and to assess its benefits we must understand bequests.

If individuals are altruistic, it is conceivable that bequests could hedge life-span risk relatively well if the bequest motive is strong. With utility deriving solely from consumption rather than from other aspects of living, then the disutility of early death could in theory be balanced by increased utility among survivors under altruistic bequests. Similarly, the additional utility deriving from late death could be offset by the impact of diminished bequests. But it is difficult to see why such fully altruistic individuals with utility only from consumption would care about any moment of life span, including the average. Probably 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. Since a bequest motive is likely to reduce both simultaneously, my results should not be highly sensitive to this assumption.

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 life-span uncertainty for reasonable parameter values.

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. Lifetime expected utility is the sum of period utilities drawn from consumption, $u(c(t))$, weighted by the force of time discounting, $e^{-\delta t}$, and the probability that the individual is alive, $\ell(t)$:

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

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 $q(t) = \log[\ell(t)] - \log[\ell(t + 1)]$.⁷ 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.

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.⁸ 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

⁷This 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.

⁸Kalemlı-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.

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.⁹

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. Other things equal, consumption will tend to be flat through age. But without annuities, $D = 1$, and the exponentially increasing mortality rate, $q(t)$, pulls 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.

⁹In 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.

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.

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,

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

for some constant K . 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, a reasonable assumption that I later relax in numerical

simulations.¹⁰

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

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

where T is a random variable, the realization of life span. I have removed the survivorship weights and instead taken the expectation over T . For a given distribution of T , we could use equation (6) to solve the integral and then find $c(0)$ by taking the expectation.¹¹ 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):

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

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

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

where K is the constant utility of being alive, and

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

¹⁰The mean life span, M , affects $c(0)$ in obvious ways. 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.

¹¹Without annuities, the presence of $q(t)$ in the consumption function precludes analytical solutions because mortality increases exponentially with age.

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.¹²

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 $T \sim N(M, S^2)$.¹³ When life spans are normally distributed, 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.¹⁴

¹²Constant absolute risk aversion over life span seems consistent with time trends in life expectancy and in S_{10} during the modern era, just as constant relative risk aversion over consumption risk matches the patterns in aggregate consumption. Average life spans have grown linearly over time, with a fairly constant standard deviation in levels, while consumption grows exponentially with steady percentage deviations.

¹³As shown by Figure 1, adult life spans are technically not normal, with both leftward skewness and leptokurtosis, or peakedness with fat tails, an indicator of different subgroup variances. Below, I show that numerical simulations show that normality reduces the cost of S as long as the discount rate is positive, so this assumption produces an underestimate of the true cost.

¹⁴It is a standard observation that period utility, which is the marginal utility of being alive in that period, should be nonnegative when modeling dynamics of life span (Rosen, 1988; Hall and Jones, 2004; Becker, Philipson and Soares, 2005). If it were negative or zero, a utility maximizing individual would choose to die. Becker, Philipson and Soares (2005) calibrate the additive utility shifter $K < 0$. This does reduce technically the cost of S through the second piece of (11), but numerical simulations confirm this effect to be small and uninteresting.

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.¹⁵ 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:

$$\bar{p}_S = -\hat{\delta}S. \quad (14)$$

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 $\hat{\delta}$, 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

¹⁵The 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 $\hat{\delta}$.

strongly with the riskiness of returns.¹⁶

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. As I will discuss later, 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 this $p_S = -\hat{\delta}S$ 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 assert that this cost seems large. Later, I will provide some context for assessing the cost relative to levels of population health across time and space. But before discussing the implications of my results, I check their robustness using numerical methods.

3.3 Numerical solutions of the full model

I can examine the sensitivity of the analytical result, $p_S = -\hat{\delta}S$, to alternative assumptions about wealth annuitization and mortality by using numerical methods. I set parameters

¹⁶According 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.

to match those used by Becker, Philipson and Soares (2005): $r = \delta = 0.03$, $\gamma = 0.8$, and $K = -16.2$. I set initial wealth at \$800,000, which is consistent with 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.¹⁷ 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. Later I reintroduce infant and young-age mortality. 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.

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 , which is 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 numer-

¹⁷I 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 (2006) topcodes age at 110, and there are few documented individuals who have survived to that age. When life spans are normally distributed with means around age 80, densities past age 110 are not miniscule. Truncating at age 110 actually creates significant skewness in the distribution, and probably changes the mean and variance. Such skewed distributions actually produce a p_S locus that better resembles that produced by realistic survivorship, probably because real life spans are skew-left.

ical 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 $r = 0.03$, or below 0.01 in the figure, $\hat{\delta}$ 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.

We 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 are 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 less interesting than one might expect. 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 intuitively tied to its other role, as the inverse of the elasticity of intertemporal substitution.¹⁸ 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 is little gain to intertemporal redistribution, and high or low interest rates do not greatly affect $\hat{\delta} \approx \delta$.

3.3.2 Realistic adult life spans

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. Colloquially, we might say that an increase in variance when there is leftward skewness reduces survivorship at young adult as well as adult ages, or across two age groups, while increasing it only at old ages. Since an individual with positive time discounting values the present more than the future, we expect

¹⁸An 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).

skewness to 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 $S_{10} = 15.66$.¹⁹

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 $p_S = -0.7$ 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.

3.3.3 Fully realistic life spans with infant mortality

Some human life spans end at birth or shortly thereafter, of course. The variance in life span attributable to infant mortality must be quite costly, but I have not yet focused on it explicitly. This is because modeling infant mortality is in a sense complex and also easy.

¹⁹I 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 discussed in the text.

The causes of infant mortality may be related to maternal health but in general are vastly different than the causes of adult death, so we should treat them differently. It is also clear that a reduction in infant mortality has more obvious benefits than a reduction in adult life-span variance. 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 expected utility of life.

The presence of infant mortality should affect p_S two ways, by reducing the average length and raising the total variability of life. By making mean life scarcer, infant mortality may lower p_S because the marginal utility of M is higher. But we have also seen that the marginal disutility of S rises with S . A rise in infant mortality and total variance could have the same effect, which would then tend to increase p_S .

To assess the cost of adult variance under fully realistic survivorship with infant mortality, I treat infant mortality as a completely separate dynamic. I fix infant deaths at their relative probability in 1999 and reestimate the compensating change in mean adult life span that offsets the reduction in S from 15.66 to 15.05.²⁰

Figure 6 depicts p_S as functions of δ when $r = 0.03$. Survivorship is now fully realistic, with infant mortality in addition to leftward skewness and leptokurtosis in the adult hump. Including infant mortality does not significantly change the qualitative results, but in the numerical models p_S is generally smaller in magnitude than it was in Figure 5 with realistic adult mortality but no infant mortality. Including infant mortality apparently increases the marginal utility of mean life more than it raises the marginal disutility of a standard deviation in life span, so p_S is less negative than before. Overall, however, the analytical solution still remains a conservative estimate of the true cost.²¹

²⁰I 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.

²¹To assess how very high rates of infant mortality may 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.

4 Discussion

4.1 The role of the discount rate

A key element is the discount rate, δ , which in this model 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, which would clearly be problematic but also seems far-fetched. Public health campaigns certainly suggest that most risky behaviors are socially undesirable, if ultimately saying little about their net private returns. A complete normative assessment thus remains elusive in this scenario, a point I concede but leave to future efforts to resolve.

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.²²

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, which the most commonly cited population health statistic, we would also like to know the level of variance faced by actual cohorts.

Differences between cohort and period S_{10} do not appear to be large, as shown by Figure

²²The less fortunate also bear a heavier burden if they have disproportionately less access to annuities, which we have found to offset perhaps one third of the cost of life-span uncertainty. This characteristic may be observationally linked to high δ , since one reason why low-SES individuals might appear to have high δ , myopia, or insufficient saving, is if they face liquidity constraints or incomplete markets. We would expect that access to annuities markets would also be poor for liquidity constrained individuals.

7. I plot both measures for the U.S. since 1900 using annual period life tables from the Human Mortality Database (2006) 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 $S_{10} = 21$.²³

4.3 Population health over time and space

The implications of equation (14) for assessing population health are significant along several interrelated dimensions. First, we can propose a new method of comparing health across industrialized nations today that simply augments the standard technique of measuring life expectancy to account for the cost of aggregate health inequality. Second, we gain a new perspective on the gaps in well-being between rich and poor countries today. Third, we have new insights into the nature and timing of demographic and epidemiological transitions that have occurred and are proceeding today in various parts of the world. And fourth, we can reassess U.S. progress against mortality over long periods of time, when declines in life-span variance were as prominent as increases in life expectancy.

In 1999, individuals in the U.S. experienced a standard deviation in life spans conditional on survival to age 10, S_{10} , equal to about 15 years. At that level, each year of standard deviation is worth about $p_S = -0.45$ year of mean life span in this model, assuming $r = \delta = 0.03$, the standard value in calibration exercises using U.S. data. In Sweden that same year, S_{10} was about 13. According to this model, individuals in the U.S. would be willing to give

²³Technically, 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.

up almost 0.9 year in mean life span to have the lower S_{10} of their Swedish counterparts.²⁴ Although 0.9 year sounds small relative to mean life span in either country, it is large when compared to the difference in means between the two countries. The mean life span conditional on survival to age 10, M_{10} , was 77.7 years in the U.S. and 80.0 in Sweden in 1999. If we account for differences in S_{10} , the total difference in population health between the U.S. and Sweden is more like 3.2 life years per person rather than 2.3, an increase of more than a third.

Another perspective suggested by these results is that the benefits of rapid gains in period life expectancy among developing countries during the past 50 years have probably been tempered by less progress against variance in life span. Valuing gains in life expectancy in addition to gains in income per capita, Becker, Philipson and Soares (2005) reveal considerable growth in total well-being, or “full income” which monetizes the value of the average length of life, among developing countries and thus much world-wide convergence. But owing to data constraints, they can only account for the changes in the first moment of life span. Although the amount of life-span variance in developing countries is unknown, it is surely higher than in industrialized countries and probably still relatively high by historical standards.²⁵ A separate research project aims to compare what limited historical data exist to recent life tables estimated by Lopez et al. (2001) and others, who account for the considerable impact of HIV/AIDS on the mean and the variance.

It seems likely that the basic convergence result of Becker, Philipson and Soares, although weakened, will remain even after accounting for life-span variance. While many developing countries have felt the burden of HIV/AIDS and have not completed their epidemiological transitions, industrialized countries have made little progress against variance since 1960 (Edwards and Tuljapurkar, 2005). Still, poor countries are probably still considerably poorer

²⁴For 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 $M_1 - M_2 = \hat{\delta}(S_1^2 - S_2^2)/2$.

²⁵High-quality data in the Human Mortality Database (2006) show that S_{10} is currently between 16 and 18 in three former Soviet republics including Russia.

in terms of a “full income” concept that accounts for the costs of variance.

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 variance in life spans. The second stage of the transition is characterized by a shift in focus toward treating chronic degenerative diseases that afflict the elderly, which tends to lengthen the average adult life span. My results suggest that was probably the optimal sequence.

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.²⁶

My findings suggest it is also worthwhile to reassess the economic value of historical gains against mortality, which we can now decompose into gains against variance and 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. Both papers reveal that the value of health improvements is very large, rivaling the value of GDP. But Nordhaus’s method places no value on reductions in variance beyond the value of increases in life years weighted by age-specific survivorship probabilities. As his results using several alternative methods suggest, this is equivalent to

²⁶To be sure, the story also involves marginal costs, since the socially optimal allocation of resources occurs where the ratio of marginal utilities of 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.

valuing gains in life expectancy at birth, which is a survivorship-weighted average. That is, he implicitly assumes agents are risk-neutral over life years.

Murphy and Topel add in a modern schedule of age-specific valuations of life years, whose shape will reflect the cost of variance. But they fix this schedule in time, which prices variance at current levels. As a result, their estimates partially capture but will underestimate the value of reduced variance, since modern levels of variance are low. Still, their results are useful in that they highlight the wide dispersion of gains across ages prior to 1950, when reductions in variance were more important.

We see exactly that pattern when we chart mortality decline during the 20th century in the U.S. The top three rows in Table 1 list life expectancy at birth, e_0 , S_{10} , the standard deviation in life spans above age 10, and the share of deaths occurring before age 10. The middle rows list the assumed discount rate, δ , and the price of S_{10} in terms of the mean, p_S , which depends on its level. The bottom rows translate S_{10} into mean life year equivalents, weight by survivorship to age 10, and subtract the cost of variance from e_0 . The last row reveals that accounting for gains against S_{10} revises upward our measure of total life-years gained by 24 percent over the 20th century, with the vast majority of the revision coming before 1950.

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.

My results imply that worldwide, health inequality must be larger between rich and poor countries than is implied by life expectancy alone, since life-span uncertainty is surely higher in developing countries. But we currently have scant evidence regarding this issue, owing to significant data problems. Estimating current life tables in developing countries is an area of much interest (Murray et al., 2003; Hill and Choi, 2004). Assessing historical trends is harder. In this country, where data is more readily available, accounting for the value of decreases in S_{10} in the during the 20th century adds about 25 percent to estimates of the total value of mortality declines.

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.

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.

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. 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.
- 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. 2007. "The Cost of Cyclical Mortality." Unpublished manuscript, April.
- Edwards, Ryan D. and Shripad Tuljapurkar. 2005. "Inequality in Life Spans and a New Perspective on Mortality Convergence Across Industrialized Countries." *Population and Development Review* 31(4):645–675.
- 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.
- 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.
- 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.
- Hill, Kenneth and Yoonjoung Choi. 2004. "The Adult Mortality in Developing Countries Project: Substantive Findings." Paper prepared for Adult Mortality in Developing Countries Workshop.
- 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. 2006. University of California, Berkeley (USA) and Max Planck Institute for Demographic Research (Germany). Available at www.mortality.org.
- Hurd, Michael D. 1987. "Savings of the Elderly and Desired Bequests." *American Economic Review* 77(3):298–312.
- Hurd, Michael D. 1989. "Mortality Risks and Bequests." *Econometrica* 57(4):779–813.
- 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.
- Kalemli-Ozcan, Sebnem, Harl E. Ryder and David N. Weil. 2000. "Mortality decline, human capital investment, and economic growth." *Journal of Development Economics* 62:1–23.
- 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.
- 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.
- 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.
- Miyamoto, John. M. and Stephen A. Eraker. 1985. "Parameter Estimates for a QALY Utility Model." *Medical Decision Making* 5(2):191–213.
- Murphy, Kevin M. and Robert H. Topel. 2006. "The Value of Health and Longevity." *Journal of Political Economy* 114(5):871–904.

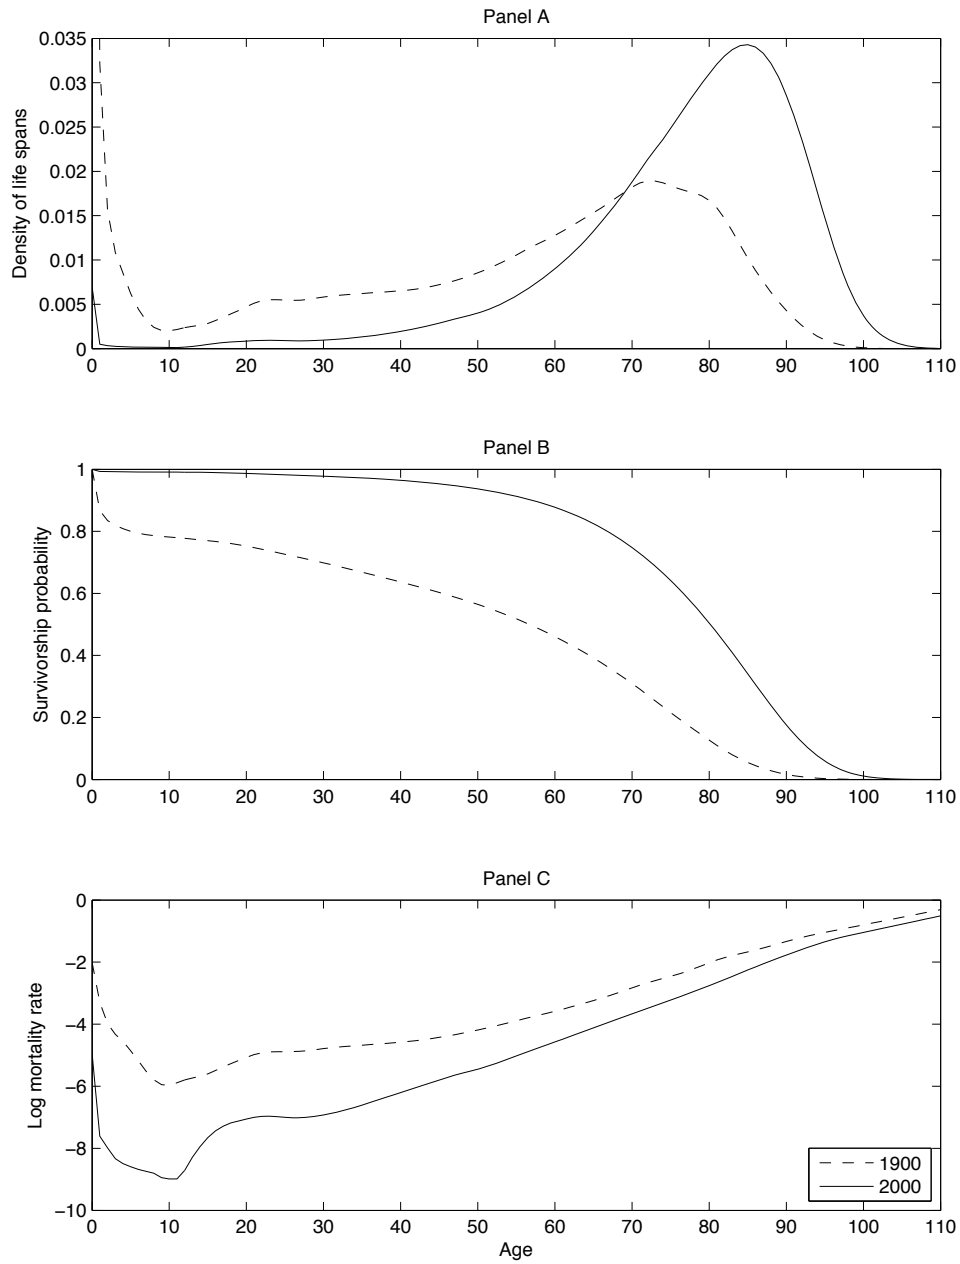
- Murray, C. J. L., B. D. Ferguson, A. D. Lopez, M. Guillot, J. A. Salomon and O. Ahmad. 2003. "Modified Logit Life Table System: Principles, Empirical Validation, and Application." *Population Studies* 57(2):165–182.
- Nordhaus, William D. 2003. The Health of Nations: The Contribution of Improved Health 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 HYE — 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.
- 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.
- 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.
- 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.

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

	Levels			Changes		
	1900	1950	2000	1900- 2000	1900- 1950	1950- 2000
e_0 , life expectancy	47.7	68.4	76.7	29.0	20.7	8.3
S_{10} , adult life-span variance	24.0	16.0	14.9	-9.1	-8.0	-1.1
deaths before age 10	21.8%	3.7%	0.9%	-20.9%	-18.1%	-2.8%
δ , the discount rate	0.03	0.03	0.03			
$p_S = -\delta S_{10}$	-0.72	-0.48	-0.45			
$p_S S_{10}$, the life-year cost of S	17.3	7.7	6.7	-10.6	-9.6	-1.0
$p_S S_{10}$ weighted by 1-deaths	13.5	7.4	6.6	-6.9	-6.1	-0.8
$e_0 - p_S S_{10}$	34.2	61.0	70.1	35.9	26.8	9.1
$\Delta(e_0 - p_S S_{10})/\Delta e_0$				1.24	1.29	1.10


Notes: Demographic data are simple averages of sex-specific period life tables presented by Bell and Miller (2005). S_{10} is the standard deviation of period life spans (life table deaths) above age 10.

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 (2006).

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 (2006). Life-table deaths at ages under 10 are set to equal deaths at age 10, and the entire distribution is rescaled to sum to unity.

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



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

Figure 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).