The Cost of Cyclical Mortality

Ryan D. Edwards*

March 19, 2008

Abstract

Sustained growth in both incomes and life spans are the hallmarks of modern development. Fluctuations around trend in the former, or business cycles, have been a traditional focus in macroeconomics, while similar cyclical patterns in mortality are also interesting and are now increasingly studied. In this paper, I assess the welfare implications of cyclical fluctuations in mortality using a standard model of intertemporal preferences. Mirroring the classic result of Lucas (1987) regarding business cycles, my findings suggest that short-term fluctuations in mortality are not very costly. Secular improvements in life expectancy and gains against static health inequalities appear to be much more important.

JEL: E32; I10; J17;

KEY WORDS: Health inequality; Uncertainty; Population health; Business cycles; Mortality

To paraphrase and adapt a famous quote by Lucas (1988), once we start to think about the long-term secular increases in human longevity that have accompanied modern growth, it is hard to think about anything else. Period life expectancy, or the average life span for a representative individual who

*Assistant Professor of Economics, Queens College, City University of New York. Mailing Address: Powdermaker 300, 65-30 Kissena Blvd., Flushing, NY 11367. `redwards@qc.cuny.edu`. I thank David Canning, Shripad Tuljapurkar, Ronald Lee, seminar participants at Queens College, CUNY, and at the Harvard Initiative for Global Health, and two anonymous referees for helpful comments. This research was partially supported by NICHD grant T32 HD 07329.

lives his or her entire life in a period, has increased at an average annual clip of about 0.2 year across industrialized countries since 1955 (White, 2002; Hall and Jones, 2007), before which progress was often even more rapid. The highest recorded life expectancy, a measure of best practices, has followed a remarkably linear trend for much longer (Oeppen and Vaupel, 2002).

As with economic growth itself, the underlying causes of these monumental improvements in well-being remain elusive. Improvements in income and nutrition, public health initiatives, and especially knowledge and technology are all potential contributing factors (Deaton, 2004), but untangling a definitive story is difficult since causality runs in multiple directions (Smith, 1999; Bloom and Canning, 2000; Cutler, Deaton and Lleras-Muney, 2006). Still, one would probably expect to see income and health trending in the same direction, since more income can buy more health, and more health allows one to produce more income. But without adequate public health initiatives, economic development can also bring deterioration in health (Szreter, 1997).

Against this backdrop of monumental improvements in life span over the long run, we can also identify sizeable short-run fluctuations in mortality that look like business cycles. The top panel in Figure 1 overlays plots of the log age and sex-adjusted aggregate mortality rate, the log of real GDP per capita, and their trends estimated using the Hodrick and Prescott (1997) filter. The bottom panel replaces log mortality with period life expectancy at birth. All three series oscillate around their long-run trends.[1] Especially past 1950, fluctuations in log mortality appear larger than those in life expectancy because the former is a simple average of age-specific mortality rates while the latter is a nonlinear geometric average.

Aside from the near simultaneity of the World War I economic boom and the massive spike in mortality due to the Spanish influenza outbreak, it is difficult to see any connection between the fluctuations in log mortality and in real income, which are trending in different directions in the top panel of

[1]The time-series properties of these data are the focus of much research, and I offer no new insights here. Unit roots are often posited in both cases; Lee and Carter (1992) propose a model for log age-specific mortality rates with a unit root, while Nelson and Plosser (1982) are unable to reject a unit root in U.S. national income data. In this paper, I first provide a visual analysis in Figure 1 of the series and their trends as estimated by the Hodrick and Prescott (1997) filter. In my empirical analyses, I examine first differences in log mortality and log income because it is analytically convenient and standard in the literature. Tapia Granados (2005) discusses in greater detail the time-series properties of mortality and GDP in the context of exploring the links between them.

Figure 1. It is somewhat easier to see short-run correlation in the lower panel, where life expectancy and income are both increasing. The faint picture that emerges is rather unexpected: life expectancy seems to fall below its trend when GDP rises above its. That is, the long-run positive relationship between life expectancy and income appears to be reversed in the short run. Lower-frequency swings in the trend lines themselves also support this view, with hourglass-shaped gaps appearing between them.

This perspective, that life expectancy appears to be countercyclical and mortality procyclical, is clearly at odds not only with the long-run relationship but also with the traditional perspective on mortality in the short run (Brenner, 1971, 1975, 1979, 2005). But as revealed by Ruhm (2000, 2003, 2007, 2008), Neumayer (2004), Tapia Granados (2005), Gerdtham and Ruhm (2006), Edwards (2008*b*) and others, an expanding body of evidence suggests that macroeconomic good times, or business cycle expansions, seem to be bad for population health. Traffic accidents appear to be strongly procyclical, but so do cardiovascular disease and other stress-related ailments, suggesting relatively broad incidence.

To be sure, these results do not necessarily contradict the finding that job loss is harmful to the health and well-being of individuals who are laid off. It could be that reductions in job stress and risk taking among the employed during a recession could produce a positive net effect on population health even though the jobless minority are negatively affected (Catalano and Bellows, 2005). Still, it is remarkable that the traditional normative perspective on fluctuations vis-à-vis the common good should be effectively turned inside-out. I began this paper by paraphrasing Lucas; can it be that we have arrived at Keynes, but only to turn him on his head, not once but twice? In the long run, we are all in fact increasingly alive; in the short run, and moreover during good times, we may be dead!

Work continues in this subfield, but a critical question that has remained unanswered is how costly are procyclical fluctuations in mortality, or more generally all cyclical fluctuations in mortality? The answer is important for guiding research and for informing policy in the same way that the cost of business cycles is an important parameter.

In this paper, I estimate the economic cost of cyclical mortality, paralleling the classic accounting by Lucas (1987) of the cost of business cycles. Although the theoretical structure is somewhat different, my results are qualitatively and quantitatively very similar to those of Lucas. The welfare cost of cyclical mortality, whether tied to business cycles or not, appears to be extremely small, both in an absolute sense and relative to either the benefits of continuing increases in average life expectancy or to the cost of static uncertainty in life spans, which could also be termed inequality in length of life. This is because fluctuations around the upward trend in life expectancy, as shown in Figure 1, are both small and by definition evenly distributed on either side of the trend. The average cost of higher mortality during an expansion is not only relatively limited, it is also partially offset by the average benefit of lower mortality during a recession.

This is not to say that we should ignore the plight of vulnerable groups who may disproportionately bear the cost. Nor should it imply that increasing life expectancy is necessarily preferable to all other pursuits, which may include reducing costly health inequalities. My results simply suggest that temporary fluctuations in average mortality, while interesting, are probably not a key public health priority.

In the rest of the paper I build my case for pricing cyclical mortality. I begin by reviewing estimates of the size and shape of fluctuations in mortality rates. I discuss the difficulty in conceptualizing the cost of an uncertain mortality rate, and then I present a method of translating it into the cost of uncertain life span. Then, drawing on related work that examines the latter (Edwards, 2008*a*), I propose a means of pricing cyclical mortality for a representative individual, and I recover an estimate that is both astonishingly small and also quite consistent with that of Lucas (1987, 2003). Finally, I discuss the implications of my results.

Quantifying cyclical mortality

As suggested by Figure 1, one could stochastically model either log mortality or period life expectancy rather interchangeably. The standard practice in the literature on procyclical mortality is to model mortality rates, in part because they are simpler to measure among population subgroups defined by age, sex, or other characteristic. Forecasts often focus on mortality rates as well, since demographers have typically viewed the dominant temporal trend as proportional decline in mortality rates (Lee and Carter, 1992; White, 2002).

Lee and Carter (1992) propose modeling log age-specific mortality rates as a random walk with drift, a specification that captures over 90 percent of the variation. I fit a simplified version of their model to the post-1946 mortality data shown in the top panel of Figure 1. My choice of sample period is motivated by increased stability in both mortality and macroeconomic variables. Ordinary least squares estimation reveals

$$\Delta \log m_t = \underset{(0.0024)}{-0.0107} + \underset{(0.0182)}{\eta_t}\ , \tag{1}$$

where m_t is the age and sex-adjusted mortality rate, and standard errors are in parentheses. It is straightforward to introduce a stationary covariate in order to explore how mortality responds to the business cycle. The most common choice in the procyclical mortality literature is the level or change in the unemployment rate. Here, I use the change in log GDP per capita, which is also standard:

$$\Delta \log m_t = \underset{(0.0030)}{-0.0159} + \underset{(0.0961)}{0.2607}\, \Delta \log GDP_t + \underset{(0.0172)}{\epsilon_t}\ . \tag{2}$$

The coefficient on $\Delta \log GDP$, call it $\gamma = 0.2607$, fits neatly into the range of estimates reported by Tapia Granados (2005). Faster growth in GDP of one percentage point is associated with about a quarter percentage point slowdown in mortality decline, which has averaged 1.07 percent each year per equation (1). As for GDP growth, a simple analogue of equation (1) reveals

$$\Delta \log GDP_t = \underset{(0.0064)}{0.0201} + \underset{(0.0241)}{\nu_t}\ . \tag{3}$$

In equation (2) there are two sources of fluctuations in the rate of decline in mortality: ϵ_t and $\gamma \cdot \nu_t$, the shock to the growth rate of GDP that is transmitted to the growth rate of mortality. These parameter estimates imply that procyclical mortality, the component associated with macroeconomic fluctuations, is responsible for a standard deviation of $\gamma \cdot \sigma_\nu = 0.0063$, while other sources of cyclical mortality comprise $\sigma_\epsilon = 0.0172$, or a level almost three times larger.

A theoretical valuation of cyclical mortality

Translating mortality into length of life

Results so far have revealed the degree of uncertainty in mortality rates, and we wish to gauge its welfare cost. Economic theory allows us to price uncertainty in rates of return; for example, by using the Consumption CAPM or related models. But there is no extant theory regarding how to price uncertainty in rates of death.

I argue that this requires translating uncertainty in mortality rates into uncertainty in the length of life. This is because standard economic models fail to capture risk aversion over mortality rates, while they can appropriately model risk aversion over length of life. The standard model of expected lifetime utility is

$$EU = E\left[\int_0^T e^{-\delta t} e^{-m(t)} u(c(t))\ dt\right], \tag{4}$$

where $m(t)$ is the mortality rate, a random variable. A mean-preserving spread in a particular $m(t)$ actually raises expected utility in equation (4) through Jensen's Inequality because $e^{-m(t)}$ is convex.

Instead of modeling variance in the mortality rate, it is helpful to focus on life span, and to decompose uncertainty in length of life into two conceptually distinct components. We can define *life-table uncertainty* as the inherent spread in length of life around its mean that arises when mortality rates rise gradually through age.[2] A convenient visualization of life-table uncertainty is the probability density function of life spans or life-table deaths in a particular year. Figure 2 shows this distribution for the U.S. in 2000. The mean, which was about 77 that year, is also known as period life expectancy at birth, which is shown over time in the lower panel of Figure 1. Smooth increases in the mean over time have been fully consistent with a fixed amount of life-table uncertainty as measured by the standard deviation of life span when the latter is primarily old-age mortality (Edwards and Tuljapurkar, 2005).

In contrast, *temporal uncertainty* in life span can be defined as the volatility in the mean length of life around its long-run trend, which is visible as the oscillations in Figure 1. This is a residual category that encompasses the very same concept of cyclical mortality we have already considered, but reinterpreted in a convenient fashion that will become clear shortly.

This decomposition is useful because we can treat both components of uncertainty in length of life as isomorphic, even though they are conceptually distinct. The two panels in Figure 3 reveal this insight graphically using actual data for the U.S. in 2000. The solid line in the top panel depicts the

[2]Demographers know this uncertainty either as a non-rectangular survival curve, a positive and finite Gompertz slope in log mortality, or a bell-shaped life-table death distribution. All are interchangeable definitions.

age schedule of log mortality, which to a first approximation is a Gompertz (1825) curve:

$$\log m_{x,t} = \alpha + \beta x, \tag{5}$$

where x is age, t is time, and α and β are constants.[3] Mortality rises gradually with age through the β parameter, which equals 0.087 in 2000. Equation (5) thus incorporates life-table but not temporal uncertainty in life span.

If the cyclical shocks to the first difference in log mortality from equations (2) and (3), $\epsilon_t + \gamma \cdot \nu_t$, are constant across age,[4] then we can rewrite equation (5) to include temporal uncertainty as well. The level of log mortality must be subject to a shock with half the variance of the shock that impinges the difference:

$$\log m_{x,t} \;\; = \;\; \alpha + \beta x + \frac{1}{\sqrt{2}} \left(\epsilon_t + \gamma \cdot \nu_t \right). \tag{6}$$

It is easy to see that the cyclical disturbance term, $\epsilon_t + \gamma \cdot \nu_t$, simply shifts the entire mortality schedule left or right. This dynamic is shown by the dashed schedules on either side of the solid line in the top panel of Figure 3. A more subtle insight is that this additive translation is also like resetting age, x:

$$\log m_{x,t} \;\; = \;\; \alpha + \beta \left(x + \frac{\epsilon_t + \gamma \cdot \nu_t}{\beta \sqrt{2}} \right). \tag{7}$$

[3] Mortality at very young and very old ages clearly does not follow this schedule, but total deaths are few at both extremes. For ease of exposition, I assume α and β do not change over time, but my results do not depend on this rather unrealistic simplification. A good fit of temporal trends in mortality can be obtained by allowing the Gompertz intercept to decrease linearly, $\alpha_t = \bar{\alpha} - gt$ for some constant g, while fixing the slope at $\beta_t = \bar{\beta}$ and thus fixing the adult life-table uncertainty (Tuljapurkar and Edwards, 2008). Under these circumstances, equation (5) implies that $\Delta \log m_x = -g$, which is analogous to equation (1).

[4] Although this is an approximation, it is a relatively good one. Tapia Granados (2005) shows there are some differences in the incidence of procyclical mortality across groups defined by age, race, and sex, and Edwards (2008*b*) reveals that the disabled appear to suffer countercyclical mortality, for example. Edwards is unable to uncover many significant differences by socioeconomic status. But as Ruhm (2008) and others have remarked, the phenomenon is remarkably wide-ranging in scope, affecting young and old similarly. Since the bulk of mortality occurs during late working age and retirement, two broad age groups that are similarly affected by procyclical mortality, the assumption of uniform incidence is convenient and appears to be relatively benign.

Redefining age in this way changes only the mean of the life table death distribution while leaving centered moments unaffected. This is shown by the addition of dashed lines in the the bottom panel of Figure 3, which plots the corresponding distributions of life table deaths by age in 2000.[5]

The upshot is that we can reinterpret cyclical mortality as additional life-table uncertainty in length of life. This is because more uncertainty in the mean length of life is approximately equivalent to more variance in the distribution around a known mean. When they are uncorrelated, life-table uncertainty and short-run temporal uncertainty are just additive layerings of the same uncertainty around life span.[6] Once we know how to price the former, we can price the latter and the total.

To summarize, suppose that in the absence of temporal uncertainty, the length of life, now T rather than x, is normally distributed with mean μ_T and variance σ_T^2, where σ_T^2 represents only the life-table uncertainty in T. As revealed by equation (7), we can interpret temporal uncertainty as shifting the mean age μ_T by $(\epsilon_t + \gamma \cdot \nu_t)/(\beta\sqrt{2})$. That is a normally distributed disturbance that we can reinterpret as additional variance around μ_T. Then if life-table uncertainty and temporal uncertainty are independent, the resulting distribution of length of life T becomes

$$T \sim N\left(\mu_T, \sigma_T^2 + \frac{1}{2\beta^2}\sigma_\epsilon^2 + \frac{\gamma^2}{2\beta^2}\sigma_\nu^2\right). \tag{8}$$

Valuing uncertainty in length of life

Now that I have translated cyclical mortality into additional life-span uncertainty, what remains is to place a value on the latter. Edwards (2008*a*) proposes a new method of pricing life-span uncertainty using a standard

[5] In both panels, I have set an arbitrarily large vertical distance between adjacent schedules at $\epsilon_t + \gamma \cdot \nu_t = 0.435$, for expository purposes. A 43.5 percent difference in mortality rates would be very large indeed, given that the average annual rate of mortality decline is about 1 percent. As shown, the Gompertz slope, β, allows us to convert that 0.435 into a horizontal difference of 5 years on either side of the original schedule in the upper panel.

[6] To be sure, independence is not a prerequisite if the covariance structure were known. Visually, there appears to be little relationship between fluctuations in life-table uncertainty, as reported by Edwards and Tuljapurkar (2005), and fluctuations in mortality decline. Their preferred measure of the former is S_{10}, the standard deviation in the age at death above age 10. A simple OLS regression of the change in S_{10} on the change in life expectancy, e_0, reveals no significant relationship between them (t-statistic of 0.04).

model of intertemporal choice in economics. That paper focuses on the consequences of large differences in static life-table uncertainty across countries at points in time and within countries over long periods of time, while this paper examines the cost of short-run temporal fluctuations. Here I provide a brief overview of the method; the details are discussed by Edwards (2008*a*).

The problem requires finding the marginal disutility of uncertainty in length of life, which in a standard intertemporal model like equation (4), primarily depends on the rate of time discounting. The discount rate is the coefficient of absolute risk aversion in length of life (Edwards, 2008*a*).[7] The intuition behind this result is that a mean-preserving spread in life span exchanges utility earlier in time, which is less heavily discounted and thus more valuable, for more heavily discounted utility later. As Edwards (2008*a*) shows, even full annuitization of wealth cannot fully hedge against life-span risk.

An analytical representation of the cost of uncertain life span can be derived when wealth is fully annuitized, markets are complete, there is no other source of risk, utility is isoelastic, and life span is normally distributed[8] with mean M and variance S^2. Under those conditions, maximization of equation (4) subject to a standard budget constraint implies that (indirect) expected lifetime utility is

$$EU \approx \mathcal{C}\,\frac{1}{\delta}\left[1 - e^{-\delta M + \delta^2 S^2/2}\right], \tag{9}$$

where $\mathcal{C}$ is a constant that depends on lifetime wealth. The price of a standard deviation in life span, S, in terms of the mean, M, is the marginal rate of substitution between them:

$$p_S = \frac{\partial EU/\partial S}{\partial EU/\partial M} = -\delta S. \tag{10}$$

The cost of an additional year in standard deviation is equal to the discount rate times the current level of a standard deviation in life span. Additional

[7]Bommier (2006) considers a more general specification of preferences with an independent parameter governing risk aversion over length of life. Experimental data are rare, and no calibration is provided. It remains to be seen whether the empirical degree of risk aversion over life span is greater or less than that implied by time discounting alone.

[8]The distribution of human life span is skew-left and leptokurtic, as shown in Figures 2 and 3. Under fully realistic mortality, which also includes a spike at birth and infancy, the cost of uncertain life span is somewhat higher than the p_S given in equation (10). This is because skewness and the infant spike remove more utility earlier in life than when length of life is normally distributed. But equation (10) remains a close approximation.

life-table uncertainty is costlier when the time discount rate is higher or when there is more uncertainty. When combined with earlier results, equation (10) reveals the cost of cyclical mortality insofar as it contributes to S, the standard deviation in length of life, which is shown in equation (8).

Calibrating the cost of cyclical mortality

As shown by Edwards and Tuljapurkar (2005), total uncertainty in adult life span in the U.S. has oscillated around $S = 15$ years since 1960. This is considerably higher than uncertainty in low-variance countries like Sweden and Japan, where currently $S = 13$ after gradual but continuous declines during the same period.

A standard estimate of the time discount rate is $\delta = 0.03$ (Becker, Philipson and Soares, 2005). Given that, equation (10) implies that the two-year difference in S we see between the U.S. and other industrialized countries is worth about one year in M. We can convert this to dollars using a price of \$200,000 per life year, which is roughly the average estimate according to Tolley, Kenkel and Fabian (1994) after updating for inflation. These results imply that static differences between countries in the level of life-table uncertainty surrounding length of life are relatively large.

The decomposition of the total variance in length of life, S^2 in equation (8), combined with the calibration results from equations (2) and (3) reveal that temporal uncertainty is small relative to life-table uncertainty. When $\sigma_T^2 = 15$, $\beta = 0.087$, $\sigma_\epsilon = 0.0172$, $\gamma = 0.2607$, and $\sigma_\nu = 0.0241$, we find that equation (8) implies

$$\begin{aligned} \text{Var}[T] &= \sigma_T^2 + \frac{1}{2\beta^2}\sigma_\epsilon^2 + \frac{\gamma^2}{2\beta^2}\sigma_\nu^2 \qquad (11) \\ &= 15^2 + 0.0195 + 0.0026 \\ &= (15.0007)^2 \end{aligned}$$

Temporal uncertainty, whether associated with the business cycle or not, adds only about 0.0007 year to the standard deviation in length of life, S. Taken alone, procyclical mortality adds only about 0.0001.

Because temporal uncertainty contributes so very little to S, the cost of cyclical mortality according to this framework is extremely low. At \$200,000 per life year, the cost of an additional 0.0007 year in standard deviation comes out to just \$70. The procyclical mortality piece on the far right of equation (11) is responsible for only \$10 of the \$70.

Discussion

This paper has shown that a standard model of intertemporal utility maximization attaches a vanishingly small average welfare cost to cyclical mortality, perhaps $70 per person each year. This result closely mirrors those of Lucas (1987, 2003), who considers the welfare cost of consumption fluctuations during business cycles. Using a representative agent model, Lucas recovers an estimate of about one-twentieth of 1 percent of consumption, or about $20 per person.

Compared to these negligible costs, the gains accruing from secular growth in incomes and life spans are overwhelmingly larger. Annual growth in consumption has averaged 2 percent. The value of annual gains in longevity, which have averaged 0.2 year, is roughly equal to annual gains in consumption (Nordhaus, 2003).

Both Lucas (1987, 2003) and I use essentially the same analytical approach, so the similarity in outcomes is not surprising. Both papers assume a representative agent with access to complete markets. These individuals display only a moderate level of risk aversion, either over the fluctuations in consumption considered by Lucas, or over the fluctuations in life span that are of primary interest here. In either case, these preferences are arguably appropriate for the average consumer. But we know that in reality, there is considerable heterogeneity in preferences and markets are incomplete.

An additional caveat that applies here is that the true degree of risk aversion over periods of life could be different than the time discount rate. This is a topic that is only beginning to be explored, and a more complete answer awaits future inquiry. We know from financial economics that attitudes toward risk and time preference are often more complicated than standard economic models suggest. Using financial market data, Alvarez and Jermann (2004) estimate a much higher cost of business cycles than Lucas did, basically because of the equity risk premium puzzle.

While there is some empirical evidence that individuals are risk averse over years of life (Edwards, 2008*a*), there is no clear consensus in health economics about whether people perceive short and long-term risks to their physical well-being as economists think they should, or about how to model preferences over length of life. Bommier (2006, 2007) prefers a baseline model with greater risk aversion than implied by the standard framework of exponential time discounting that I use, but he provides no empirical support of this hypothesis. In his framework, agents shift consumption earlier in life in order to hedge against uncertain life span even if they are fully annuitized. When time discounting is exponential, as I assume, fully annuitized agents do not shift consumption in response to mortality risk, but they are still hurt by life-span uncertainty. Without more data on actual preferences over length of life, it is difficult to assess the empirical degree of risk aversion, and much work remains to be done in this subfield. In the meantime, it seems reasonable to proceed with standard modeling techniques and calibration settings, as I have here.

Differential incidence of cyclical mortality is a potentially critical issue that I have assumed away. This is because the extant literature currently points in no clear direction; the phenomenon of procyclical mortality seems to be quite broadly based. But it remains an open question whether some groups are disproportionately more at risk of poor health and death during fluctuations. We know that different groups face different levels of life-table uncertainty, with lower socioeconomic status correlated with shorter as well as more uncertain length of life (Edwards and Tuljapurkar, 2005). If there were differential incidence in cyclical mortality that favored the well-off, my results could change.

Another element I do not directly consider is the possibility that procyclical mortality could actually hedge consumption risk. Since only a small portion of all temporal uncertainty in mortality appears to be attributable to business cycle fluctuations, the hedging benefit of procyclical mortality is not likely to outweigh the negative impacts of fluctuating mortality. Still, it could prove interesting to revisit this issue directly by modeling uncertainty in both life span and consumption simultaneously.

My results certainly do not diminish the catastrophic nature of the 1918 influenza pandemic or other temporary adverse shocks to mortality. Lucas's insights did not deny the staggering cost of the Great Depression. But my results do suggest that as a society, it appears that we should focus our energies on reducing static health inequalities and maintaining overall progress against mortality rather than weathering the vicissitudes. In any event, as discussed by Edwards (2005) and Catalano and Bellows (2005), it is not entirely clear what the policy options would be in the case of procyclical mortality. Would we ever trigger a recession in order to improve public health?

References

Alvarez, Fernando and Urban J. Jermann. 2004. "Using Asset Prices to Measure the Cost of Business Cycles." *Journal of Political Economy* 112(6):1223–1256.

Backus, David K. and Patrick J. Kehoe. 1992. "International Evidence on the Historical Properties of Business Cycles." *American Economic Review* 82(4):864–888.

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.

Bloom, David E. and David Canning. 2000. "The Health and Wealth of Nations." *Science* 287(5456):1207–1209.

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

Bommier, Antoine. 2007. "Mortality Decline and Aggregate Savings." Unpublished working paper, May.

Brenner, M. Harvey. 1971. "Economic Changes and Heart Disease Mortality." *American Journal of Public Health* 61(3):606–611.

Brenner, M. Harvey. 1975. "Trends in Alcohol Consumption and Associated Illnesses: Some Effects of Economic Changes." *American Journal of Public Health* 65(12):1975.

Brenner, M. Harvey. 1979. "Mortality and the National Economy." *The Lancet* 314(8142):568–573.

Brenner, M. Harvey. 2005. "Commentary: Economic Growth is the Basis of Mortality Rate Decline in the 20th Century— Experience of the United States 1901–2000." *International Journal of Epidemiology* 34(6):1214–1221.

Catalano, Ralph and Benjamin Bellows. 2005. "Commentary: If Economic Expansion Threatens Public Health, Should Epidemiologists Recommend Recession?" *International Journal of Epidemiology* 34(6):1212–1213.

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

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

Edwards, Ryan D. 2005. "Commentary: Work, Well-Being, and a New Calling for Countercyclical Policy." *International Journal of Epidemiology* 34(6):1222–1225.

Edwards, Ryan D. 2008*a*. "The Cost of Uncertain Life Span." Unpublished manuscript, March.

Edwards, Ryan D. 2008*b*. "Who Is Hurt by Procyclical Mortality?" Unpublished manuscript, February.

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.

Gerdtham, Ulf-G. and Christopher J. Ruhm. 2006. "Deaths Rise in Good Economic Times: Evidence from the OECD." *Economics & Human Biology* 4:298–316.

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.

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.

Hodrick, Robert J. and Edward C. Prescott. 1997. "Postwar U.S. Business Cycles: An Empirical Investigation." *Journal of Monetary, Credit and Banking* 29(1):1–16.

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

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

Lucas, Robert E. 1988. "On the Mechanics of Economic Development." *Journal of Monetary Economics* 22(1):3–42.

Lucas, Robert E. 2003. "Macroeconomic Priorities." *American Economic Review* 93(1):1–14.

Nelson, Charles R. and Charles I. Plosser. 1982. "Trends and Random Walks in Macroeconomic Time Series: Some Evidence and Implications." *Journal of Monetary Economics* 10:139–162.

Neumayer, Eric. 2004. "Recessions Lower (Some) Mortality Rates: Evidence from Germany." *Social Science & Medicine* 58(6):1037–1047.

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.

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

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

Ruhm, Christopher J. 2007. "A Healthy Economy Can Break Your Heart." *Demography* 44(4):829–848.

Ruhm, Christopher J. 2008. Macroeconomic Conditions, Health and Mortality. In *Making Americans Healthier: Social and Economic Policy as Health Policy*, ed. Robert F. Schoeni, James S. House, George A. Kaplan and Harold Pollack. New York: Russell Sage.

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

Szreter, Simon. 1997. “Economic Growth, Disruption, Deprivation, Disease, and Death: On the Importance of the Politics of Public Health for Development.” *Population and Development Review* 23(4):693–728.

Tapia Granados, José A. 2005. “Increasing Mortality During the Expansions of the US Economy, 1900–1996.” *International Journal of Epidemiology* 34(6):1194–1202.

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

Tuljapurkar, Shripad and Ryan D. Edwards. 2008. “Variance in Death and Mortality Decline.” Unpublished manuscript prepared for the 2008 Annual Meeting of the Population Association of America.

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

Figure 1: Trends in log mortality, life expectancy, and log real income per capita in the U.S. since 1900

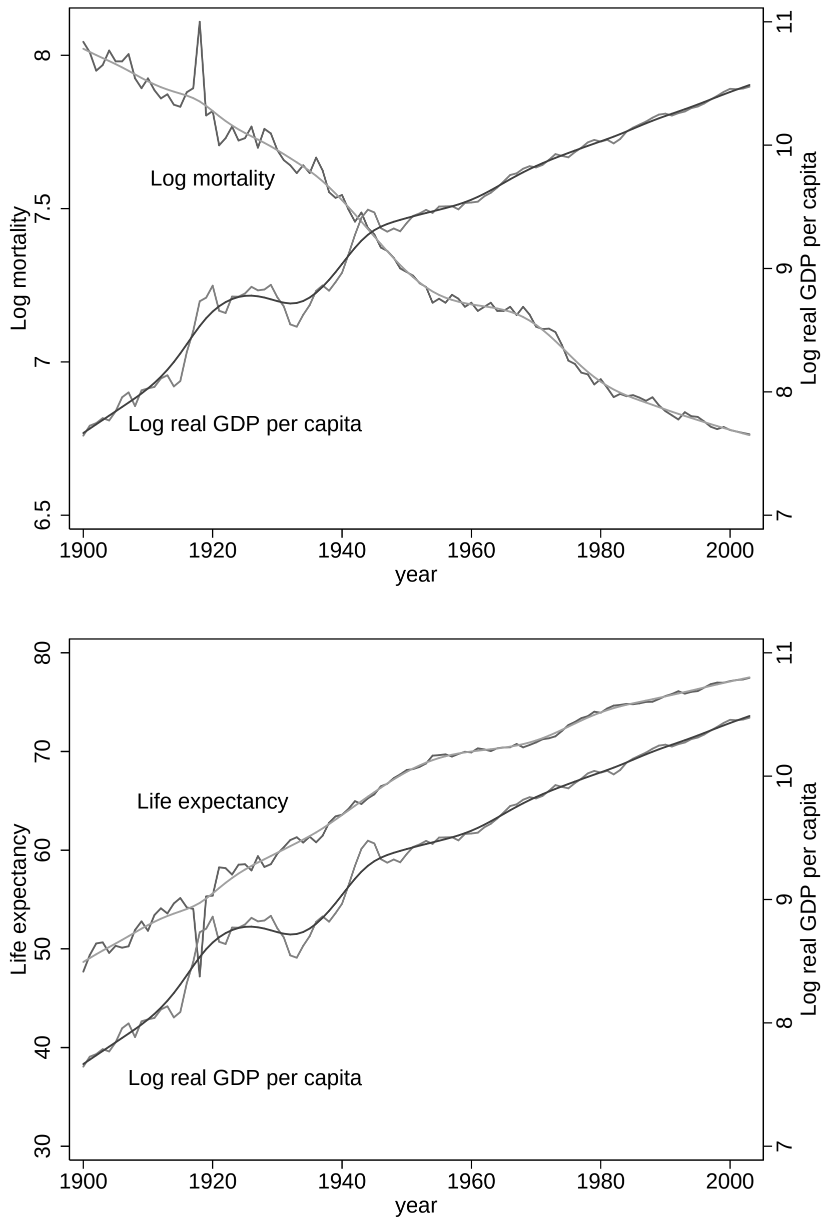

Notes: Data on log age and sex-adjusted mortality are taken from the *Historical Statistics of the United States*, series B-167-180, and the 1997 and 2004 Social Security Trustees Reports. Data on real GDP per capita are provided by the Bureau of Economic Analysis and the *Historical Statistics*. Life expectancies are averages by sex based on data in Bell and Miller (2005). Trends are calculated using the Hodrick and Prescott (1997) filter with smoothing parameter $\gamma = 100$, per Backus and Kehoe (1992).

Figure 2: The probability distribution of human life span in the U.S. in 2000

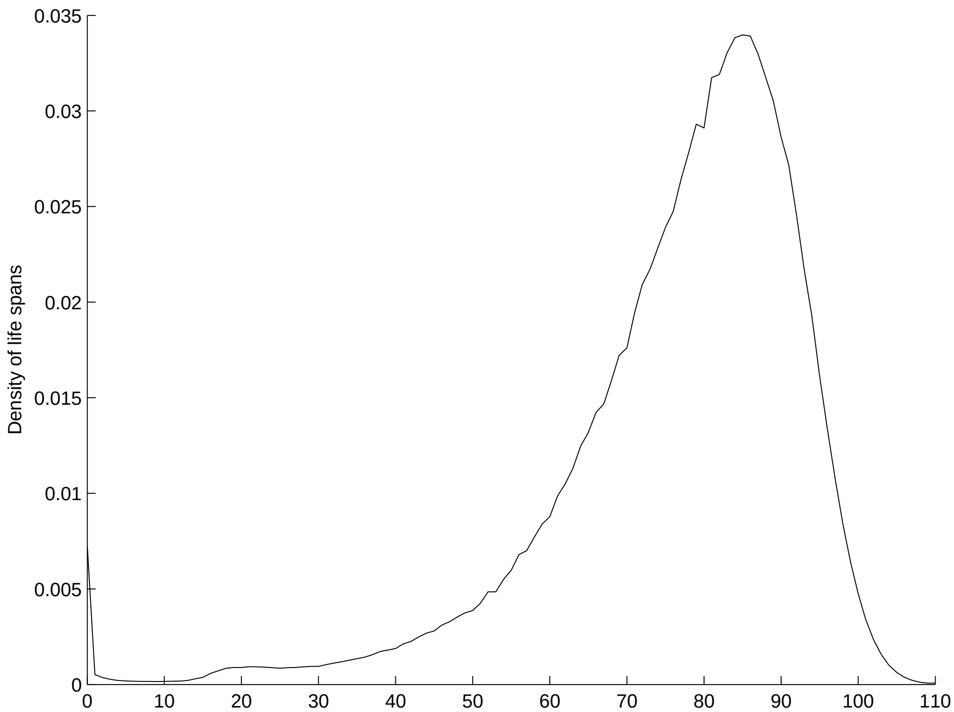

Notes: Data are from Bell and Miller (2005) and are the average by sex of the life table deaths.

Figure 3: Additive shocks to Gompertz mortality are additive shocks to mean life span

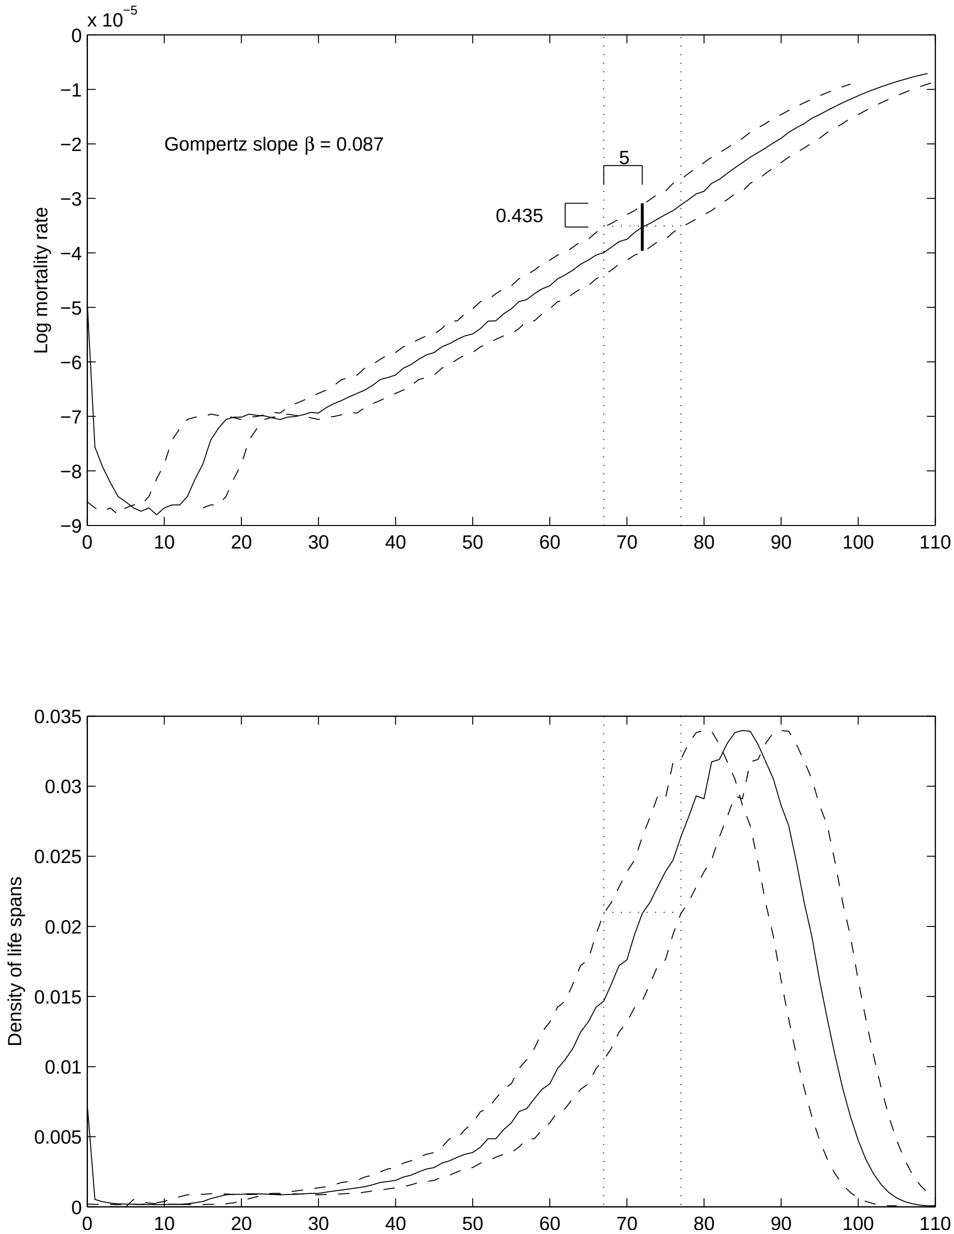

Notes: Data are from the 2000 U.S. period life table reported by Bell and Miller (2005) and are the average of sex-specific quantities. See the text for details. The upper panel plots log age-specific mortality rates for 2000 (solid line) separated by additive shocks of 0.435 and −0.435 from two fictitious schedules. The lower panel plots the life-span (life table deaths) distributions corresponding to all three, where the Gompertz slope in the upper panel has translated an additive shock to log mortality into an additive translation of the distribution of life spans.