

Inequality in Life Spans and a New Perspective on Mortality Convergence Across Industrialized Countries

Ryan D. Edwards and Shripad Tuljapurkar*

September 27, 2005

Abstract

The second half of the twentieth century witnessed much convergence in life expectancy around the world. We examine differences in the age pattern of mortality between countries over time to show that inequality in adult life spans, which we measure with the standard deviation of life table ages at death above age 10, S_{10} , is increasingly responsible for the remaining divergence in mortality. We report striking differences in level and trend of S_{10} across industrialized countries since 1960, which cannot be explained by aggregate socioeconomic inequality or differential external-cause mortality. Rather, S_{10} reflects both within and between-group inequalities in life spans and conveys new information about their combined magnitudes and trends. These findings suggest that the challenge for health policies in this century is to reduce inequality, not just lengthen life.

*PRELIMINARY. PLEASE DO NOT CITE OR QUOTE WITHOUT PERMISSION. Edwards: Postdoctoral Scholar, RAND Corporation, redwards@rand.org. Tuljapurkar: Morrison Professor of Population Studies, Department of Biological Sciences, Stanford University. The authors wish to thank Ronald Lee, seminar participants at RAND, and several anonymous referees for helpful comments. Edwards acknowledges support from the Morrison Institute and from NIA grant T32 AG000244. Tuljapurkar acknowledges support from NIA grant 1P01 AG022500. This paper uses data supplied by the National Heart, Lung, and Blood Institute, NIH, DHHS from the National Longitudinal Mortality Study. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the National Heart, Lung, and Blood Institute, the Bureau of the Census, or the National Center for Health Statistics.

The human condition has improved tremendously during the course of modern development. At the beginning of the nineteenth century, life expectancy at birth, e_0 , hovered between 25 to 40 years (Maddison, 2001). Industrialization and unprecedented growth in per-capita incomes coincided with significant gains in e_0 , which by 1960 reached roughly 70 years among industrialized countries and since then has increased at about 0.2 year of age for every year of time (White, 2002). Since the early nineteenth century, record female life expectancy increased at an annual rate of almost 0.25 years (Oeppen and Vaupel, 2002).

These monumental gains in e_0 have been fairly equally distributed across international boundaries. White (2002) reports convergence in e_0 among 21 industrialized countries during the postwar period. Developing countries have also displayed convergence toward industrialized levels since 1950, as described by Wilson (2001). But life expectancy is only a summary measure of the age pattern of mortality, although it is an important one. A larger research question is whether the age pattern of mortality has also shown convergence between advanced countries over time. To address this issue, we examine full distributions of life table ages at death, of which life expectancy is the average, or first moment.

We use a relative entropy metric, the Kullback-Leibler divergence, to compare changes over time in the distributions of life table ages at death among industrialized countries, using combined-sex data. While we find evidence of convergence, we identify cases of lingering differences in mortality between otherwise similar countries. After decomposing our metric of convergence into portions attributable to the mean and variance of age at death, we reveal that current differences in the age pattern of mortality are attributable to differences in variances in adult age at death.

We measure adult life span variability by S_{10} , the standard deviation in the life table age at death past age 10. We find starkly heterogeneous trends in the life span variance S_{10} across advanced countries since 1960, in contrast to the convergence in average age at death among this same group. Viewed relative to long-term historical experience with S_{10} , largely dominated by long periods of relative stagnation interrupted by a vast reduction in variance during the epidemiological transition, these modern-day differences in S_{10} take on added meaning. While gains in life expectancy have been remarkably steady both overall and across international boundaries, gains against life span variance have been more scarce. Thus our findings on S_{10} naturally lead us to expand traditional perspectives on mortality convergence

with several important caveats. But they also raise a host of new questions to which we then turn.

First, we specifically motivate our focus on S_{10} as opposed to the many alternative measures of dispersion or inequality that are available. We explain how S_{10} is the most practically useful and intuitively revealing measure of dispersion in the current context. It falls naturally from our decomposition of the convergence metric, and it is easy to interpret relative to the mean. We relate our findings to earlier work on inequality in life table ages at death, and to work on health disparities.

Next, we ask whether cross-country differences in S_{10} based on combined-sex data are in fact driven by sex-specific differences. The course of modern development has brought mortality gains to men and women at different times based on a complex interplay of biological differences, social behavior and norms, gender equality, and scientific advancement. Many cross-country differences in life expectancy that exist today can be traced to sex-specific factors such as differential alcoholism prevalence, smoking, and access to medical technologies, among other things. Here we find no evidence that systematic heterogeneity in S_{10} across countries is closely related to differences in S_{10} by sex.

For a single country, S_{10} places an upper bound on the sum total of life span inequality between citizens, which is the sum of between-group inequalities and complex within-group inequalities. Thus S_{10} cumulates well-known racial and socioeconomic inequalities in mortality such as those famously explored by Kitagawa and Hauser (1973). So it is natural to ask whether countries with high S_{10} , notably the U.S., have high variance mainly because of high between-group inequality. The answer turns on the relative magnitude of within-group variance to between-group variance. Thus we next explore how S_{10} is related to socioeconomic inequality within the U.S. Examining U.S. data, we use several decompositions the literature suggests are responsible for large between-group variation (by race, education, income, and external versus other causes of death); in each decomposition, we compare within-group and between-group variance in age at adult death. Excluding mortality from external causes, such as accidental and inflicted deaths, smooths fluctuations somewhat but does not make S_{10} in the U.S. look like that of other countries in either level or trend. Racial heterogeneity and socioeconomic inequality also do not appear to explain divergence in life span variance, although there are clear racial and socioeconomic divides in S_{10} . We show that for many such partitions of the population, within-group

variance substantially outweighs between-group variance in all cases. In fact, within-group variance is so important that we observe higher socioeconomic strata within the U.S. experiencing higher S_{10} than citizens of other countries.

Returning to cross-country differences, our next step is exploring how aggregate socioeconomic inequality relates to S_{10} . We find clear evidence that patterns in S_{10} across rich countries are not explained by inequality in income or in education. Rather, S_{10} captures a key inequality that is not measured by other aggregate statistics and has no simple relationship with socioeconomic status. Thus our findings reveal a novel perspective on mortality inequality that bears implications for conceptualizing and modeling mortality across socioeconomic groups, and for better understanding the state of convergence in mortality across countries. In addition, S_{10} reflects an important uncertainty that likely impacts a wide range of economic decisions made over the life cycle, the efficient functioning of insurance markets, and government finances.

In the sections that follow, we elaborate on each of these points. First, we propose an explicit measure of the relative importance of life span means and variances in adult mortality convergence, and then we demonstrate the increasingly important role of variance. Next, we directly analyze adult life span variance as measured by S_{10} among a subset of seven representative industrialized countries. We also compare trends since 1960 with historical developments in variance. We discuss whether sex differences in mortality play a role in the heterogeneity we see in S_{10} across countries. We relate S_{10} to alternative measures of variance in health and life span, and we motivate our focus on it in particular. Then we present several different decompositions of S_{10} , first using aggregate data broken down by race and cause of death, and then with more detailed data on individuals in the U.S. National Longitudinal Mortality Study. We also depict the convoluted relationships between S_{10} and traditional measures of aggregate socioeconomic inequality. These suggest that S_{10} is a complex function of socioeconomic status and that S_{10} contains information beyond that in the SES variables considered here. Finally, we discuss the implications of our findings.

Convergence in distributions of ages at death

Cross-country convergence in mortality is typically measured using life expectancy, the life-table mean age at death. Life expectancy is a convenient

summary measure of mortality, more informative than an age-standardized mortality rate and simpler than a Gompertz slope. But the full distribution of life-table ages at death is characterized not only by the mean, but also by the variance and other moments that describe the age pattern of deaths. While the mean is the most important moment, mortality convergence is better measured using the entire distribution of ages at death.¹

A graphical comparison quickly suggests why this is so. Figure 1 depicts distributions of life-table ages at death for Sweden and the United States in 1999, using data by single years of age for both sexes combined supplied by the Human Mortality Database (2004), hereafter HMD. Infant mortality is represented by spikes on the left-hand side of the graph, while adult mortality can be described as bell curves around old-age modes. Despite the fact that modal ages at death were almost identical in the two countries, life expectancies at birth were actually quite different, at 76.8 years in the U.S. and 79.5 years in Sweden. The variance in adult ages at death, or the horizontal width of the old-age hump, was larger in the U.S. than in Sweden. Both distributions show negative skewness, and both are leptokurtic, or highly peaked and fat-tailed. Dispersion around the mean age at death is clearly important in the overall similarity of the two mortality schedules.

How can we formalize differences between the curves in Figure 1 in order to help us understand their sources? Information theory provides us with tools to analyze directly the similarity of two distributions and decompose their differences. In particular, the Kullback-Leibler divergence, hereafter KLD, is a useful statistic that is related to a log likelihood ratio and Akaike's Information Criterion.² A larger KLD represents an increasingly dissimilar pair of distributions, and minimizing the KLD is similar to maximizing likelihood or optimizing the fit of a model to reality.

Before we calculate the KLD, we truncate all distributions below age 10 so as to remove effects of the infant mortality spike. We do this for two reasons. While trends in infant mortality are important, even a small level of infant mortality has a huge effect on the moments of the distribution of ages at death. We focus on the ages where most death now occurs and which are the focus of current interest in the health and mortality of aging populations. If we leave out infant mortality, we find that the full KLD between any two age distributions of death can be accurately decomposed into a piece that reflects differences between the mean ages of adult death and a piece incorporating differences between the variances in ages at adult death. This approximation requires the assumption of near-normality in ages at death, which, as we have

seen in Figure 1, is more reasonable if infant mortality is excluded.³

We calculate the KLD for each country in each year relative to a static “best practices” distribution, that of Sweden in 2002. This is done in order to account for temporal as well as cross-sectional differences in mortality. If instead the baseline distribution were chosen contemporaneously, for example comparing Denmark in 1980 with Sweden in 1980, temporal trends, which most demographers would agree are of prime importance in describing mortality decline, would effectively be ignored.⁴ The exact interpretation of the level of the KLD varies according to application; here, $KLD = 0.1$ corresponds to a difference in mean ages at death of about 6 years, while $KLD = 0.01$ represents a difference of about 2 years.⁵

Figure 2 shows convergence in adult mortality as measured by the KLD for seven countries since 1960.⁶ The overall picture is one of near-monotonic convergence over time toward Sweden’s mortality schedule in 2002. Japan’s record of convergence is particularly noteworthy as the most rapid, with slightly increased divergence after 1990 due to its surpassing Sweden as the leader in life expectancy. Denmark, meanwhile, converged the slowest and remained the most divergent by the end of the period. The U.S. lingers above the other five nations at higher levels of divergence.

As was suggested by Figure 1, trends in variance turn out to explain a nontrivial part of the lingering divergence evident in Figure 2. Using our approximation method, we decompose contributions to the KLD into portions attributable to only the mean age at adult death and to only the variance. Figure 3 depicts the results of this decomposition for Canada, Denmark, France, and the U.S. The solid lines are the approximate total KLD relative to Sweden’s age-at-death distribution in 2002, the dashed lines are the portion attributable to differences in means, and the dotted lines show the partial effects of different variances. Although differences in means tend to dominate during most of the period for most countries, divergent variances gain in importance for France and the U.S. For those two countries as well as for Canada, variances actually surpass means in importance after 1990.

Although there has indeed been much convergence in mortality among industrialized countries since the Second World War, we have presented evidence in this section of some lingering divergence in adult mortality. A more revealing discovery is that the source of this divergence appears increasingly to be an heretofore underappreciated component of mortality: the variance around the adult mean age at death. How has life span variance evolved in advanced countries over the past half-century, what explains those patterns,

and what are their implications?

Variance in adult life span, S_{10}

Measures of variation in life span depend crucially on whether infant and child mortality is included or excluded, and we argue that new and valuable insights are gained from focusing only on adult variance. Wilmoth and Horiuchi (1999) and Shkolnikov, Andreev and Begun (2003) provide many insights into the variation of life-table ages at death over the entire age range using an array of measures, including interquartile range, standard deviations, and Gini coefficients. Along with Kannisto (2000), their ultimate focus is the compression of human mortality and evidence of rectangularization around an upper limit to life span, rather than mortality convergence across countries. We find that unconditional variance, measured over the entire age range, is a poor measure for informing analysis of mortality convergence, because uniform declines in infant mortality obscure the more significant (for total deaths) dynamics in adult mortality. In a recent study, Cheung et al. (2005) investigate rectangularization and mortality decline by decomposing the survival curve into infant, premature, and “normal” deaths at or after the mode age at death. This leads them to focus on the dispersion in deaths at ages beyond the mode, and thus on death rates at the highest ages; our analysis of variance contributed by deaths at all adult ages provides quite different insights.

Unconditional standard deviations in ages at death, S_0 , have declined quite homogeneously across industrialized countries since 1960, as exhibited by Figure 4. Levels of S_0 may be different across countries, but trends are broadly similar. A graph of log infant mortality rates during the same period reveals almost identical trends (not shown). That is, S_0 is heavily influenced by declines in infant and child mortality and thus offers few insights into the dynamics of the old-age hump.

Once infant mortality is removed, by conditioning on survival until age 1, 5, or 10 and then calculating standard deviations, a radically different picture emerges.⁷ Figure 5 graphs our preferred measure, S_{10} , the standard deviation in ages at death over age 10, for the same seven countries. Now, not only are levels of S_{10} different, trends are quite different as well. The U.S. suffers the highest S_{10} , around 15–16 years, with little relief in sight given its general lack of trend. France’s experience echoes that of the U.S.,

but at a slightly lower level closer to 15 years. Canada, like the U.S. a high-immigration country, had similar S_{10} up until 1980, when it began to trend lower, toward Danish levels of around 14 years. Denmark started the period like Sweden, which experienced low and stable S_{10} , but then lost ground and did not recover. Japan experienced rapid gains early and converged to the low levels enjoyed by the United Kingdom and Sweden by 1975. Altogether, the gap between high- S_{10} and low- S_{10} countries has averaged 2.5 years of life since 1960, a spread of nearly 20 percent around a typical S_{10} of 14.

Relative to some historical periods, recent declines in S_{10} have been quite moderate. While each year's record female life expectancy has registered steady and remarkably linear increases since 1840 (Oeppen and Vaupel, 2002), best-practices variance in the age of adult death has followed a more complicated path. As shown by Figure 6, the lowest S_{10} recorded each year among 14 high-income countries in the Human Mortality Database for both sexes combined was relatively stagnant prior to about 1875, when it entered a period of rapid decline lasting roughly 75 years.⁸ We find that progress against variance was unprecedentedly rapid during the mortality revolution of the late 19th and early 20th centuries, echoing the results of Wilmoth and Horiuchi (1999), whose preferred measure of variance is the interquartile range over all ages. But after about 1950, S_{10} entered a period of relative stagnation. The static differences in levels and trend in S_{10} across countries that we see in Figure 5 are arguably more meaningful when viewed relative to the post-1950 stagnation of the frontier apparent in Figure 6. That is, gains against adult life span variance are more dear when they are atypical for the time.

There is a vast literature on dispersion in mortality and health, and a pertinent question is how S_{10} relates to previously researched concepts. Wilmoth and Horiuchi (1999) and Shkolnikov, Andreev and Begun (2003) explore a variety of measures of life span variability, including the interquartile range, the Gini index, and the standard deviation, among many others. The message we take from these efforts, which are centered around different research questions than ours, is that there are some tradeoffs between the various measures, but that many are qualitatively indistinguishable. We prefer S_{10} because it falls naturally from our examination of the Kullback divergence, and because we believe it is a more intuitive conceptualization of life span variability than alternative measures, some of which tend to be unnecessarily arcane. Unlike income distributions, modern distributions of adult ages at death are virtually normal and show limited skewness, so the traditional

advantages of a percentile based approach do not apply.

Our emphasis on variance in mortality as opposed to the mean outcome has clear precedents in the literature on health inequalities. Murray, Frenk and Gakidou (2001) survey the diverse approaches in this area, which include examining the distribution of life expectancies measured across groups, typically geographic but also otherwise defined. We prefer to examine variance in individual life spans, measured by life table deaths, rather than variance in average life spans across groups. As a more disaggregated statistic, S_{10} is a more direct measure of health inequality between individuals. We prefer its broadness because of the additional insights it provides, which we discuss later in the paper. Gakidou, Murray and Frenk (2000) argue that differences in health expectancy across groups are more informative than differences in healthy life span across individuals because the former remove the component that is “simply due to chance.” In fact, we find later that a large part of aggregate S_{10} is within-group variation in life span, which probably reflects underlying uncertainty. Such uncertainty also appears to differ between groups, which suggests it is actually an important ingredient, rather than mere background noise, in overall inequality.

We acknowledge that many nonfatal health events are important, and thus life span is only one indicator of health status. But S_{10} should be as good a measure of health inequality as e_0 is of average population health, and the complementarity between them is clear. Given that e_0 is frequently although not exclusively used to measure population health, we feel that conceptualizing S_{10} as an indicator of health inequality is warranted, and it is especially useful in better interpreting e_0 .

What accounts for such divergent trends in S_{10} among countries that are experiencing much convergence in life expectancy? Why is S_{10} highest in the U.S.? In the next section, we explore these questions using aggregate vital statistics, a panel study of mortality at the individual level in the U.S., and international macroeconomic data.

The sources of S_{10}

Sex differences in mortality and S_{10}

We know that changing sex differences have been responsible for many other patterns in mortality that we find interesting. During recent decades, for ex-

ample, the gap in life expectancy between males and females in the U.S. has rapidly widened and then stabilized. In other countries, trends in the sex differential have followed similarly stark patterns that have frequently diverged at least temporarily from developments elsewhere. Thus it is clear that sex differentials in mortality are a prime suspect for understanding differential trends in S_{10} .

Sex differentials in mortality are well documented (Waldron, 2000) and appear to reflect a combination of biological and behavioral influences. Males typically suffer from higher in utero mortality, and hormonal differences appear to place males at greater risk of heart disease (Waldron, 1995). Sex ratios do not differ substantially across industrialized countries, however, which suggests that if constant, these basic biological differences could not explain differentials in S_{10} .

But sex differentials in mortality tend to vary considerably across genetically similar populations (Gjonça, Tomassini and Vaupel, 1999), implying that behavioral or environmental differences, which could shift S_{10} , are important. Pampel (2002) argues that part of the narrowing sex differential among industrialized countries is due to changes in female smoking behavior. It is conceivable that males may tend to engage in other types of risky behavior more often than females, perhaps also through their choice of occupations as well as through recreational behavior. Less clear is whether preferences, opportunities, or social norms regarding such behavior actually differ between industrialized societies, so the question of whether S_{10} is affected by trends in sex-specific mortality remains open.

We find that S_{10} among females is indeed lower than for both sexes combined, by about 1 year.⁹ This represents reduced life span uncertainty among women that is presumably tied to less exposure to risks, or to less socioeconomic or biological heterogeneity. But removing male mortality produces a level rather than a growth effect on S_{10} , and relative rankings between countries tend not to change much. Two noteworthy differences are France and Denmark; in France, female S_{10} is roughly 1.5 years lower than for both sexes combined, while in Denmark it is only 0.5 year lower. These results coincide with those of Pampel (2002), who reports evidence of lighter smoking among French females, heavier smoking among Danish females, and sex differentials in mortality that appear to reflect those behavioral differences. But while sex-specific trends may explain some part of S_{10} , virtually the same trends that are apparent in Figure 5 remain after removing male mortality.

External-cause mortality and S_{10}

Among the many causes of deaths, accidents, homicides, and suicides, commonly referred to as external causes of mortality, tend to impact younger adults much more than the elderly. A common perception is that violent deaths are disproportionately prevalent in the U.S., where firearms are more widely owned, automobile travel is more common, and crime rates are high. Whether S_{10} may be merely reflecting differences in external cause mortality is surely a key question. The dispersions through age of other causes of death are also interesting, but we omit an extended discussion for two reasons. First, our focus is on temporal change, and there are well-known problems with the consistency of diagnoses and of coding systems over time. We believe these are minimal in the case of external causes. Second, we are aware of persistent differences in disease prevalence across populations that appear to be the product of a complex interweaving of genetic, social, and behavioral factors (Goldman and Takahashi, 1996; Rimm and Stampfer, 2004). Thus we leave such inquiry to future efforts and limit our focus to external causes of death.

External causes of death can be conceptualized as reflecting risks that derive from the environment through local socioeconomic conditions and geography. It is conceivable that modern economic growth could heighten external risks, perhaps through expanding the use of automobiles and other forms of transportation, access to increasingly powerful firearms, or increasing returns to theft and other crimes. If this is true, some of the postwar slowdown in S_{10} might be attributable to increases in external cause mortality. If such risks were more prevalent in some countries, they might account for some of the observed heterogeneity in S_{10} . Wilmoth and Gullickson (2001), for example, show that a portion of the U.S. disadvantage in mortality among young adults can be traced to AIDS and external-cause mortality.

We removed deaths due to external causes using data on cause of death by 5-year age group from the World Health Organization Mortality Database (2004), and then we recalculated S_{10} using nonexternal mortality only.¹⁰ The net effect is a level decrease in each country's S_{10} of about 1–1.5 years in every chronological year. Thus external mortality clearly does contribute to S_{10} , but it accounts for less than a tenth of the total standard deviation. Some intertemporal smoothing of S_{10} occurs as a result of the decomposition, which suggests that transitory shocks to life span uncertainty may be tied disproportionately to external-source mortality. Overall, however, trends

in S_{10} remain unaltered. This is clear evidence against the notion that post-war stagnation in S_{10} might be associated with heightened risks of death by external causes brought on by rapid economic growth.

Relative rankings of the countries remain almost completely unaffected, and where there is change, it is not always in the expected direction. Britain's relative position in S_{10} actually deteriorates somewhat once external causes are excluded, which contravenes the notion that violent deaths might be high in the UK by European standards. The U.S. and France still display the highest adult variability, Sweden the lowest, and the interesting trends in Canada, Denmark, and Japan remain. Thus while a U.S. disadvantage in adult mortality rates may be attributable to higher rates of accidents, homicides, and suicides, the U.S. disadvantage in the variance of age at death is not. External-cause mortality does not appear to explain much of the heterogeneity in S_{10} among these countries.

Race and S_{10}

Could differences in racial composition affect S_{10} ? In the U.S., racial inequalities in adult mortality are well documented and likely reflect underlying socioeconomic differentials (Preston and Taubman, 1994; Tuljapurkar and Boe, 1998). France, second-highest in S_{10} after the U.S., is home to a substantial Muslim minority. Canada is also racially heterogeneous owing to high levels of immigration, and it too exhibited high S_{10} , at least prior to 1980.

Official policies on the collection of data by race or ethnicity vary widely across countries. For simplicity, we focus on mortality differentials in the U.S. between whites and African Americans, and we obtain life tables for those two racial categories and for all Americans from the Berkeley Mortality Database. The BMD, the precursor to the expanded HMD, presents aggregates derived from individual-level data provided by the National Center for Health Statistics.¹¹

Figure 7 plots S_{10} for whites only, African Americans, and for all races combined in the U.S., alongside total S_{10} for France and Canada. Variability among African Americans is considerably higher than among whites, between 2–3 years or around 15–20 percent higher, and there is much greater volatility in African-American S_{10} . Fluctuations in variability have been temporally synchronized for the two groups, however. This suggests uniformity over time in the aggregate sources of life span risk, which may then be differentially felt by population subgroups. The fact that aggregate S_{10} became smoother

when external causes were removed suggests crime-related deaths might be one such aggregate source that differentially affects racial groups. Infectious disease may be another such source, if race proxies for socioeconomic status or perhaps geography.

Measured among U.S. whites alone, S_{10} is still higher than aggregate S_{10} in France and far above Canadian levels, especially by the end of the period. Thus racial differentials in mortality do not explain very much of the relatively high level of S_{10} or its lack of trend in the U.S. It seems doubtful that a more complete accounting of the effects of racial differences among all seven countries would be particularly revealing.

Much of the racial differences in U.S. mortality are thought to proxy for socioeconomic differences. A natural question to consider next is whether heterogeneity in S_{10} simply reflects different levels and trends in socioeconomic inequality across advanced countries.

Socioeconomic inequality and S_{10}

There is a long tradition of work in public health and epidemiology exploring how income inequality may affect life expectancy (Wilkinson, 1992). Concrete evidence of a simple relationship is elusive, however, at least among advanced countries (Gravelle, Wildman and Sutton, 2002; Deaton, 2002; Lynch et al., 2004). Whether this is due to the complicated relationship between mortality rates and e_0 , factors other than income such as technology, or linear returns in health interventions, it would seem that average life span reflects few aspects, if any, of inequality.¹²

Our focus here is solely on dispersion, both in socioeconomic status as well as in life span, and we are interested in how variances in each may be interrelated. A causal link running from income to health would imply such a relationship, as would a link running from health back into income. Currently, there is much interest in the field of health economics regarding which direction causality runs and whether the direction changes with age (Smith, 1999; Adams et al., 2003). The partial answers that have emerged thus far suggest that causality is complicated, with evidence that it is not constant through age. We believe that as long as the relationship is always positive, which both theory and evidence tend to support, the variances in income and in health should be positively related. Thus the precise direction of causality should not matter for our purposes.

We examine the link between socioeconomic variance and S_{10} in two ways.

First, we directly decompose S_{10} by socioeconomic strata using longitudinal data on individuals in the U.S. National Longitudinal Mortality Study (NLMS). Then we explore the relationship between aggregate dispersion in income and in education versus S_{10} . We find clear evidence that socioeconomic inequality is important for variance in adult life span, but we also demonstrate that S_{10} measures inequality that is more than just socioeconomic in nature.

S_{10} among socioeconomic subgroups

The U.S. NLMS is a panel dataset of over half a million individuals who were interviewed in Current Population Surveys conducted around 1980 and then tracked for nine years (Rogot et al., 1988, 1992).¹³ Socioeconomic data were observed only at the beginning of the period, while detailed data on the time and cause of death are provided by matches to the National Death Index. Since fairly wide temporal swings in S_{10} are apparent in Figure 5, we construct a life table based only on mortality in the first year of the sample rather than several years' worth, as is more typically done with this dataset. Income, notorious for its wide year-to-year fluctuations, may be a poor proxy for true socioeconomic status, so we also separately decompose S_{10} by education, a much more stable measure for adults.

For ease of exposition, we identify two socioeconomic strata using either income or education. When we use income, we sort individuals according to whether they are in the first quintile of household income. With education, we sort according to whether individuals are high school graduates, roughly two-thirds of the sample. Then we calculate life tables for both sexes combined in each group, plot smoothed distributions of ages at death using a standard kernel density estimator, and present conditional means and variances.

Figure 8 depicts ages at death for the two income groups, while Figure 9 plots distributions by educational status. Both tell the same story: adults in lower socioeconomic strata not only suffer shorter average life spans, M_{10} , they also endure greater variability, S_{10} . The gap in average life span between individuals in the first income quintile and those in the top four-fifths is 5.5 years, while the difference in standard deviations is nearly 2.5 years. Similarly, high school graduates live an average of 5 years longer than their less educated counterparts, while enjoying a standard deviation that is 2 years lower.

These findings are significant in several regards. Historically, distributions of ages at death became strongly leptokurtic, or sharply peaked and fat-tailed, in advanced countries after the mortality transition.¹⁴ Figures 8 and 9 reveal why this might be so. If population subgroups have different variances, then the overall distribution will exhibit heteroscedasticity, or fat tails.

Secondly, the discovery of heterogeneous life span variances among subgroups calls into question how we typically model mortality. The usual assumption of proportional hazards in a Gompertz-style model of age-specific mortality rates, such as in a Cox regression, is inconsistent with different subgroup variances. To see why, consider log mortality rates that differ by a constant between groups. That specification is equivalent to simply rescaling age by a constant. An additive shift in the distribution of deaths by age changes only the mean and not the variance nor any other centered moments. Researchers who model mortality may want to consider whether variance is an important dynamic before imposing the assumption of proportional hazards. Proportional hazards cannot capture the dynamics of variance. We discuss the possible implications of life span variance for individuals, markets, and governments in the next section.

Finally, although socioeconomic differences in S_{10} are illuminating, our findings suggest that high S_{10} in the U.S. is not solely due to socioeconomic inequality. The text in Figure 8 shows that when the bottom income quintile is removed, S_{10} falls to 14.4. Based on Figure 5, this is a drop of 1 year in aggregate S_{10} , or about as large an effect as is gained by restricting to females only. Even without controlling for socioeconomic inequality in countries other than the U.S., which would surely lower the relevant baselines, this is still a relatively high level of S_{10} . Similarly, Figure 9 suggests that even if everyone in the U.S. had a high school diploma, S_{10} would remain fairly high. Neither income nor education alone seems to explain U.S. exceptionalism.

S_{10} and aggregate socioeconomic inequality

Although the paucity of large longitudinal datasets severely limits our ability to control for socioeconomic inequality in a panel of countries, we can instead examine how aggregate measures of inequality correlate with S_{10} . We recognize that there are limitations of such aggregate comparisons, however, and we will focus on relative temporal trends in order to reduce the contaminating effects of unobserved variables.¹⁵ At a minimum, the comparison can tell us whether S_{10} is a unique measure of aggregate inequality.

We combine aggregate income inequality statistics from two sources: for six of the countries, we use the high-quality Luxembourg Income Study (LIS) Microdatabase (2003); for Japan, we must use the Deininger-Squire (1996) dataset. Since incomes tend to grow multiplicatively, we use the Gini coefficient rather than the standard deviation to measure income inequality, a standard technique.¹⁶ Thomas, Wang and Fan (2000) provide Gini coefficients of educational attainment over time in five of the seven countries we consider in detail.

Figure 10 plots S_{10} against the Gini coefficient of household income for our seven countries since roughly 1970. Each country's data points are connected through time with starting and ending points labeled chronologically. Since there could very well be country fixed effects at work, we believe the temporal variation, or the slope of the lines, is more informative than the cross-sectional variation, the scattering of the points. Indeed, if one were to ignore the lines and focus on the scattered cloud of points, one might assert, mistakenly in our view, that there is an upward sloping, positive relationship between S_{10} and aggregate income inequality. But except in the case of France, where there is little change in either variable between 1979 and 1994, and possibly Canada, which does seem circuitously to follow an upward sloping path, temporal trends largely refute this view. Income inequality has generally increased while S_{10} has either remained level, seen most distinctly in the case of the UK, or declined. These longitudinal patterns provide no support for the notion that S_{10} merely reflects aggregate income inequality.¹⁷

Shifting our focus to inequality in education provides an even more dramatic exposition. Although we lose the UK and Sweden due to lack of data coverage, we discover an even more convoluted relationship between S_{10} and Ginis of educational attainment between 1960 and 1990, as revealed in Figure 11. First of all, the cross-sectional variation suggests no clear relationship at all. Even excluding the U.S., a clear outlier with very high S_{10} and very low educational inequality, observations are arrayed in a cloud. Secondly, trends over time follow completely opposite paths. Japan and Denmark move in opposite North-South directions, while the U.S. and France move in opposite East-West directions. Motion in all four directions is represented, along with counter-clockwise spiralling on the part of Canada. It is clear that aggregate inequality in educational attainment has no link at all to life span inequality.

To be sure, we could also chart e_0 against income or education and probably uncover similar unexplained trends in life expectancy. We could also examine e_0 relative to income inequality and find trends that contravene

prevailing wisdom (Deaton, 2002). Our point here is merely that S_{10} is another aggregate indicator of well-being that is basically orthogonal to other well-known measures. In particular, S_{10} measures an aggregate inequality that is not measured by income or educational inequality.

Summary: a unique and revealing measure of inequality

We have seen how our preferred measure of variance in the adult age at death, S_{10} , is indeed related to all of the usual covariates of mortality. External causes of mortality directly account for some of S_{10} and appear to amplify temporal fluctuations; males experience higher variance as well as lower mean life spans; African Americans appear to suffer permanently higher variance than white Americans; and less-educated and poorer individuals experience higher variances in addition to lower means. Although S_{10} clearly reflects known influences on mortality, we believe it is an important, new indicator for four reasons.

First, it is a convenient measure of a key aggregate health inequality. It neatly complements life expectancy, which is widely interpreted as average population health. S_{10} falls directly from our decomposition of the Kullback-Leibler divergence in mortality schedules, and it is a more intuitive concept than alternative measures of variability in life spans, while remaining qualitatively similar. It places an upper bound on the combined effects of within and between-group inequality on overall life span inequality. S_{10} also reveals progress toward reducing aggregate health inequality within a country, and it facilitates international comparisons.

Second, S_{10} is a unique measure of aggregate inequality. It has not simply followed trends over time in either income or educational inequality, and it does not reflect cross-sectional differences in either of those more commonly referenced indicators.

Third, it appears to reveal fundamental differences in levels and trends of health inequality across similar low-mortality economies. While socioeconomic status and other mortality influences can explain part of S_{10} , we have shown that they do not explain enough of it to eliminate exceptionally high S_{10} in the U.S., which is the most obvious pattern that emerges. We have also shown that other important findings, such as consistently high variance in France, increasing variance in Denmark, or sharply decreasing variance in Canada after 1980, cannot be trivially explained. What we can loosely term “background” variance or uncertainty in life span appears to vary across

advanced countries in ways that we do not yet understand.

Fourth, S_{10} reveals a new concept of differential mortality at the individual level that is not usually captured by traditional modeling techniques and that may be very meaningful. Cox proportional hazards may correctly fit differences in mean ages at death, but they cannot capture differences in variances, which are significant.

Inequality in life span is arguably the most fundamental inequality there is among human populations. But how important is S_{10} for other things that we care about? In the next section, we discuss a wide range of implications, ranging from insights into aggregate trends to understanding individual behavior.

Implications of S_{10}

Mortality convergence and forecasting

We began this paper with a discussion of the role that variance plays in international mortality convergence. Our findings deserve repeated emphasis because they contrast with those of researchers who have examined trends in e_0 .

In expanding the focus to variances in the age at death in addition to means, we facilitate a meaningful yet parsimonious decomposition of mortality convergence through the Kullback-Leibler divergence. This decomposition shows that while averages have indeed steadily converged, substantially different variances are driving international mortality differentials. When viewed in this way, the evidence is much less supportive of convergence toward a single mortality schedule.

The implications for forecasting are clear. Projection models must account for trends in the variance or risk incorrectly forecasting a crucial aspect of mortality that affects convergence. Simple models based on trends in life expectancy will obviously miss trends in variance and thus overstate future mortality convergence, but less clear is how traditional models that are in wide use today implicitly treat variance. In a future study, we plan to formally assess the structure and predictions of common mortality models in this regard.

Working, saving, and investing

Modern theories of economic behavior over the life cycle derive from the insights of Modigliani and Brumberg (1954), who posited that individuals rationally plan for key life events and presented evidence that they do. Life-cycle saving is what economists call accumulating just enough assets while working to finance retirement until death. While the simplest version of this theory is not the whole story, its basic approach remains a cornerstone of economic theory (Browning and Crossley, 2001; Dynan, Skinner and Zeldes, 2002). Thus at the individual level, we expect the length of life to be a key parameter in many decisions, although we acknowledge that some individuals may not engage in this type of forward-looking behavior.

By definition, the average length of life should have the largest impact on economic behavior. It is theoretically conceivable that higher moments might have no impact on decision making under certain circumstances. If individuals do not care about risk or are somehow able to pool it effectively, higher moments may not matter. Both possibilities seem unlikely. It is also reasonable to ask whether individuals are even aware of uncertainty surrounding the age at death. Although there clearly is heterogeneity in perceptions, we know that individuals tend to gauge their survivorship probabilities quite well on average (Hurd and McGarry, 1995, 2002). Thus while they may not understand the concept of variance as well as a statistician might, they do appear to have a fundamental grasp of the uncertainty.

We have seen that entire populations, as well as subgroups within those populations, uniformly face life span risk, which suggests that at least some component is not diversifiable. Government policies and markets help to distribute risk, and we will discuss them shortly, but for now we focus on how variance in life span may affect individual choices and well-being in several key dimensions. First, although basic schooling is compulsory in many countries, people still choose exactly how much education to acquire. Later in life, they also choose whether to continue working or retire. And as suggested by basic life-cycle theory, individuals must also choose how much to consume versus save over their life spans.

It is worth reemphasizing that the direction of causality between mortality and socioeconomic status is complicated (Smith, 1999; Deaton, 2003) and remains a topic of much research and active interest (Adams et al., 2003; Smith, 2003). With multiple pathways running in either direction, there is no simple answer to the question of whether health causes SES or SES causes

health. As we remarked earlier, evidence suggests that causality may even change direction during the life cycle, perhaps several times. Thus while we are accustomed to thinking that education, income, and wealth cause health outcomes, the interplay between these variables is more complex. There clearly is much to be gained by more work in this area. In this section, our goal is to discuss ongoing efforts that explore the economic implications and sources of life span uncertainty, rather than to advance a particular view of health-SES causality.

If individuals or their parents perceive that the life spans of their children are risky, they may be reluctant to invest in education. One way to see this is by the argument that education is a costly, irreversible, and non-tradeable investment. Individuals who perceive much life span risk may prefer the greater relative certainty of consumption to the uncertainty of investment in human capital. Another perspective, offered by Kalemli-Ozcan (2002), is in terms of the quantity-quality theory of childbearing. If parents are risk averse in their preferences for surviving children, then a decline in the uncertainty regarding their children's survival may decrease the number of children desired. When planning to have fewer children, parents can invest more in each one's education. In developed countries, of course, mortality tends to be low past infancy, which dilutes the salience of this point somewhat. But we know that mortality due to external causes is high among adolescents and early adults in some groups in the U.S. Our finding of higher S_{10} among those groups, who also typically have less education, is consistent with this story.

Many people work in order to finance retirement. Among other things, the enjoyment of retirement years depends on savings accumulated while working and being alive to enjoy them, ignoring bequests for now. Individuals who view retirement years as the reward for working may retire earlier when life spans are more uncertain, if they are unwilling to face the possibility of an abbreviated retirement. Defined benefit pension programs like Social Security, which lessen the financial risks associated with living too long, would increase this possibility. On the other hand, life span risk may reflect risks to health, which in turn may represent considerable financial risks. People who perceive significant health risks may choose to work longer in order to accumulate precautionary savings. The net impact of life span uncertainty on retirement decisions is theoretically unclear. Kalemli-Ozcan and Weil (2002) find that trends toward earlier retirement in the U.S. during the 20th century may be attributable to stagnation in the standard deviation

of the age at death.

Saving is an integral part of individual behavior over the life cycle, because it represents the transfer of resources from the present into the future. When the length of life becomes more or less uncertain, what happens to saving? Surprisingly little research has addressed this question directly, although there is much interest in life cycle saving in general. Normally, saving increases when individuals perceive increased risks to their income or to their health (Dynan, Skinner and Zeldes, 2002). But life span risk is different. People value leaving bequests, but it is not clear that they would save more if the chance of dying young rose along with the chance of dying old. Much work remains to be done in understanding saving behavior under conditions of life span uncertainty.

Markets for insurance and annuities

For markets to function efficiently, buyers and sellers need good information about what they are exchanging. In the cases of life and health insurance and annuities, it is likely that the buyers know something that the sellers of such policies do not. Individuals know their own family histories and behaviors much better than others do, so they probably have private information about how healthy they are and how long they are likely to live.¹⁸ If sellers are uncertain how long prospective buyers will live, they will price their policies higher than what the average buyer would pay, so as not to lose money if their policies are purchased only by long-lived individuals.

It is easy to see how high S_{10} may thus be extremely detrimental for private markets for insurance and annuities. We have shown how much of S_{10} in the U.S. is within-group variance. Under such circumstances, sellers cannot use group identities to differentiate between types of buyers and decrease their exposure to risk. One solution to this problem is to require all individuals to purchase insurance and annuities, perhaps through taxation and public provision. This is commonly done in advanced economies, of course, and we next discuss the implications of S_{10} for such arrangements.

Public old-age support

Social Security and other public old-age support programs effectively operate as insurance against living too long. Since they are funded by payroll taxes, such programs redistribute risk away from the elderly and toward workers.

Are there costs associated with these benefits that might be affected by the amount of life span variation?

Intuitively, variance in inputs matters for outcomes that are nonlinear functions of their inputs, by Jensen's Inequality. In the case of life span and Social Security, the relevant question is whether the total present value of lifetime benefits paid is a nonlinear function of life span. In one way, lifetime benefits are actually a decreasing function of life span. Since benefits increase only with inflation through cost of living adjustments and not interest rates, the present discounted value of a retiree's benefits is a decreasing function of time and thus of life span. Through this effect, increased variance is actually good for program finances and bad for pensioners. But survivor benefits clearly increase the costs of life span variance. When pensioners die early, they typically have young survivors, who collect many benefits, although at a reduced rate than the decedent had. The costs generated by retirees who die late are not fully offset by reduced costs for those who die early. An additional strain on the system is created by the between-group variance component of S_{10} . Low-income workers who contributed less to the system will die earlier than high-income workers, so longer-lived individuals will also draw larger pensions. As low-income pensioners die off at faster rates during the aging of each cohort, average benefits rise and raise program costs.

Conclusion

We began this paper by evaluating mortality convergence across industrialized countries since the Second World War. In our view, comparing full distributions of ages at death across countries is a natural extension to comparing life expectancies alone. We argue that it is more meaningful for assessing the true extent of convergence and for predicting future trends. But our findings not only bear important implications for how we perceive mortality convergence, they also offer new insights into relative well-being along the most critical dimension, that of life span.

Convergence in life expectancy has indeed been strong during the past 50 years. But once entire distributions of ages at death are considered, rather than just their means, key differences appear. Convergence in infant mortality, while perhaps not absolute, is strong enough to swamp estimates of life span variance that do not condition on early survival. Our preferred measure of variance, S_{10} , the standard deviation of life span above age 10, reveals stark

differences in levels and trends in adult variance among advanced countries since 1960. In some countries, variances have actually become more important than means in explaining overall mortality convergence.

In a sense, achieving “best practices” mortality increasingly means reducing inequalities as opposed to pushing the frontier. It is true that declines in S_{10} among industrialized countries have uniformly slowed since about 1950, even while increases in mean ages at death continue in dramatic fashion. But the postwar experiences of Canada and Japan demonstrate that high variance need not be a permanent condition, and that large improvements are possible. Deterioration in Denmark’s variance is a reminder of the other, more ominous possibility.

We have shown that while S_{10} reflects a variety of underlying, well-known influences on mortality, it appears to be a unique measure of a key component of human well-being. It is a composite measure of between and within-group inequality in life span that is not closely correlated with aggregate income or educational inequality. When we decomposed S_{10} using microdata from the National Longitudinal Mortality Study, we found that residual within-group inequality, or what we might term as “background” inequality or uncertainty, remained extremely important. While lower socioeconomic strata are exposed to greater life span variance as well as to shorter average life spans, higher strata in the U.S. still suffer high S_{10} relative to average citizens of other advanced countries. The sources of differential background inequality in life spans between nations remain unclear and await further research.

Finally, we have discussed the implications of our findings beyond mortality convergence and inequality in well-being. It is likely that life span variance or uncertainty, whether associated with socioeconomic status or representing background risk, is important for economic decisions made by individuals, for the efficient functioning of private insurance markets, and for government funding of old-age support programs. In this paper, we have merely outlined the relevant mechanisms that imply key roles for variance. Theoretical as well as empirical explorations of the various causal pathways are important directions for future efforts to take. We believe that achieving a deeper understanding of the impacts as well as of the sources of adult life span variance are important new frontiers in social science.

Notes

¹An alternative is to compare schedules of age-specific mortality rates across countries. In principle, mortality comparisons contain the same information as comparisons of the distribution of age at death. In practice, the age pattern of mortality hazards is not easy to interpret (Aalen and Gjessing, 2001), and for individuals the age at death is arguably more meaningful than age-specific probabilities of death.

²Given two distributions with densities $p_1(x)$ and $p_2(x)$, the Kullback and Leibler (1951) measure of the divergence of p_2 from the baseline, p_1 , is given by $KLD(p_1, p_2) = \int_{-\infty}^{\infty} p_1(x) \log(p_1(x)/p_2(x)) dx$.

³Roberts and Penny (2002) show that when p_1 and p_2 are normal densities, the KLD reduces to $K(p_1, p_2) = \log(\sigma_2/\sigma_1) + \sigma_1^2/(2\sigma_2^2) - 1/2 + (\mu_1 - \mu_2)^2/(2\sigma_2^2)$.

⁴We also examined trends in the KLD measured contemporaneously, relative either to Sweden in each year, or to the life expectancy or life span variance leader in each year. Some qualitative results change drastically when such measures are used, since as remarked they do not capture temporal change in mortality. But the key insight of this section, namely that life span variance is of increasing importance in explaining lingering divergence, remains unaltered.

⁵The KLD approximation described in note 3 reduces to $KLD = (\mu_1 - \mu_2)^2/392$ when $\sigma_1 = \sigma_2 = 14$, which is roughly the postwar average of S_{10} in our sample. Thus if the difference in mean ages at death falls by 1 percent, the KLD falls by approximately 2 percent.

⁶Here, as in the rest of the paper, we focus on Canada, Denmark, France, the United Kingdom, Japan, Sweden, and the United States. In an earlier version of this paper, we also examined Austria, Switzerland, Finland, (Western) Germany, Italy, the Netherlands, and Norway. We exclude the latter seven and focus on the first set for clarity of exposition, since the latter set does not contain any unique cases. Results for all 14 countries, chosen because they are high-income countries for whom the Human Mortality Database contains high-quality life tables, are available upon request from the authors.

⁷We calculated S_1 , S_5 , and S_{10} and found that each was significantly different from S_0 and not very different from each other. It is possible to condition on reaching any age, of course, but intertemporal comparisons of a quantity like S_{65} are not as straightforward nor as meaningful as one might think. Trends in S_{65} over time tend to be dominated by the rightward march in densities associated with rapid advances in life expectancy. This continually introduces new and large probability weights at age 65 that were previously assigned to ages 64, 63, and so on, and were thus not included in S_{65} . This dynamic tends to inflate S_{65} significantly over time, obscuring true changes in the width of the old-age hump. As long as the truncation age x is chosen so that densities in the neighborhood of x remain small during the entire period, S_x will be an unbiased measure of the width of the old-age hump.

⁸Separate trends in sex-specific S_{10} are very similar to the trend for both sexes combined. Record-low S_0 has also followed a similar historical path.

⁹For the sake of brevity, we omit graphical depictions of these and other results. They are available from the authors upon request.

¹⁰Although measures of variance in the age at death are relatively robust to the choice of

age aggregation, we prefer to use single years of age. The primary impediment to accurate measurement is top-coded ages past 85, which as shown in Figure 1, is near the mode age at death for advanced countries. We therefore use WHO data to identify the share of deaths by 5-year age group attributable to external causes, and then we apply that share to HMD data by single years of age. Since external-cause mortality primarily impacts younger age groups, the effect of age top-coding in the WHO data is minimal.

¹¹Elo (2001) and others have shown that African American mortality rates may be mismeasured either because births are underregistered or ages are misreported. We constructed alternative measures of S_{10} using Elo's corrected life tables for African Americans and found few important differences relative to the BMD measures.

¹²Infant mortality rates do appear to be positively related to aggregate income inequality Hales et al. (1999); Lynch et al. (2001). This suggests that mortality among other age groups with high dependency, such as the very old, might also be affected by income inequality.

¹³The public-use datafile does not specify exactly when individuals were interviewed for purposes of confidentiality, so the temporal precision of any mortality estimate is limited.

¹⁴We found this pattern (not depicted) when examining historical life tables provided by the Human Mortality Database (2004). Prior to 1900, distributions of adult ages at death were actually platykurtic, or flat-topped. An example of a platykurtic distribution is one that is uniform. Such trends in kurtosis seem to be consistent with the effects of an epidemiological transition away from infectious disease and toward noncommunicable causes.

¹⁵Aggregate cross-sectional variation, while informative in a descriptive sense, may not be very meaningful in any structural sense. That is, there are many reasons why we might expect different countries to experience socioeconomic inequality and S_{10} in different proportions. Chief among them are the stark differences in health care delivery systems and insurance coverage across countries. Examining trends in socioeconomic inequality and S_{10} over time helps us in this regard, but only if institutions and other factors are either fixed or changing in the same way across all countries. We are therefore cautious in interpreting what results we may find from aggregate data.

¹⁶Since the Second World War, distributions of ages at death have shifted additively, as partially evidenced by linear increases in e_0 White (2002); Oeppen and Vaupel (2002). This justifies our use of the standard deviation to measure inequality in ages at death. In contrast, Shkolnikov, Andreev and Begun (2003) propose Gini coefficients based on the distribution of ages at death. Their qualitative results, along with those of Wilmoth and Horiuchi (1999), who compared many competing measures of variance, suggest to us that our choice of the standard deviation rather than a Gini is perfectly reasonable. A standard deviation of adult ages at death may in fact be conceptually preferable. It is more intuitive than the Gini, and it complements thinking about the mean age at death. In any event, since there have been additive rather than log-additive changes over time in the distribution of ages at death, S_{10} is an appropriate measure of dispersion.

¹⁷Long-term historical trends in income inequality (not depicted) provide some additional support for this view. While S_{10} began to decline around 1875, as shown in Figure 6, income inequality within industrialized countries remained high into the 20th century and did not begin to fall until the rise of the modern welfare state during the interwar

period (Lindert, 2000; Bourguignon, 2002).

¹⁸In fact, preserving such information asymmetries under privacy laws is a central part of public policy in most advanced countries. In insurance markets, individuals may also unconsciously act as if they knew they had private information when they do not think they do. If an individual is risk averse, she may not only desire more insurance, but also end up living longer if risky behavior is detrimental to health.

References

- Aalen, Odd O. and Håkon K. Gjessing. 2001. "Understanding the Shape of the Hazard Rate: A Process Point of View." *Statistical Science* 16(1):1–22.
- Adams, Peter, Michael D. Hurd, Daniel McFadden, Angela Merrill and Tiago Ribeiro. 2003. "Healthy, wealthy, and wise? Test for direct causal paths between health and socioeconomic status." *Journal of Econometrics* 112(1):3–56.
- Berkeley Mortality Database. 2004. University of California, Berkeley. Available at www.demog.berkeley.edu/~bmd/.
- Bourguignon, François. 2002. "Inequality Among World Citizens: 1820–1992." *American Economic Review* 92(4):727–744.
- Browning, Martin and Thomas F. Crossley. 2001. "The Life-Cycle Model of Consumption and Saving." *Journal of Economic Perspectives* 15(3):3–22.
- Cheung, Siu Lan Karen, Jean-Marie Robine, Edward Jow-Ching Tu and Graziella Caselli. 2005. "Three Dimensions of the Survival Curve: Horizontalization, Verticalization, and Longevity Extension." *Demography* 42(2):243–258.
- Deaton, Angus. 2002. "Commentary: The convoluted story of international studies of inequality and health." *International Journal of Epidemiology* 31:546–549.
- Deaton, Angus. 2003. "Health, Inequality, and Economic Development." *Journal of Economic Literature* 49(1):113–158.
- Deininger, Klaus and Lyn Squire. 1996. "A New Data Set Measuring Income Inequality." *World Bank Economic Review* 10:565–591.

- Dynan, Karen E., Jonathan Skinner and Stephen P. Zeldes. 2002. "The Importance of Bequests and Life-Cycle Saving in Capital Accumulation: A New Answer." *American Economic Review* 2(92):274–278.
- Elo, Irma T. 2001. "New African American Life Tables from 1935–1940 to 1985–1990." *Demography* 38(1):97–114.
- Gakidou, Emmanuela E., Christopher J. L. Murray and Julio Frenk. 2000. "Defining and Measuring Health Inequality: An Approach Based on the Distribution of Health Expectancy." *Bulletin of the World Health Organization* 78(1):42–54.
- Gjonça, Arjan, Cecilia Tomassini and James W. Vaupel. 1999. "Male-female Differences in Mortality in the Developed World." *MIPDR Working Paper* 1999-009.
- Goldman, Noreen and Shigesato Takahashi. 1996. Old-Age Mortality in Japan: Demographic and Epidemiological Perspectives. In *Health and Mortality Among Elderly Populations*, ed. Graziella Caselli and Alan D. Lopez. New York: Oxford University Press pp. 157–181.
- Gravelle, Hugh, John Wildman and Matthew Sutton. 2002. "Income, income inequality and health: what can we learn from aggregate data?" *Social Science & Medicine* 54:577–589.
- Hales, Simon, Philippa Howden-Chapman, Clare Salmond, Alistair Woodward and Johan Mackenbach. 1999. "National infant mortality rates in relationship to gross national product and distribution of income." *The Lancet* 354(9195):2047.
- Human Mortality Database. 2004. University of California, Berkeley (USA) and Max Planck Institute for Demographic Research (Germany). Available at www.mortality.org.
- Hurd, Michael D. and Kathleen McGarry. 1995. "Evaluation of the Subjective Probabilities of Survival in the Health and Retirement Study." *Journal of Human Resources* 30(supp):s268–s292.
- Hurd, Michael D. and Kathleen McGarry. 2002. "The Predictive Validity of Subjective Probabilities of Survival." *The Economic Journal* 112(482):966–985.

- Kalemli-Ozcan, Sebnem. 2002. "Does Mortality Decline Promote Economic Growth?" *Journal of Economic Growth* 7(4):411–439.
- Kalemli-Ozcan, Sebnem and David N. Weil. 2002. "Mortality Change, the Uncertainty Effect, and Retirement." *NBER Working Paper* 8742.
- Kannisto, Väinö. 2000. "Measuring the Compression of Mortality." *Demographic Research* 3(6).
- Kitagawa, Evelyn M. and Philip M. Hauser. 1973. *Differential Mortality in the United States: A Study in Socioeconomic Epidemiology*. Cambridge, MA: Harvard University Press.
- Kullback, Solomon and Richard A. Leibler. 1951. "On Information and Sufficiency." *The Annals of Mathematical Statistics* 22(1):79–86.
- Lindert, Peter H. 2000. Three Centuries of Inequality in Britain and America. In *Handbook of Income Distribution*, ed. Anthony B. Atkinson and François Bourguignon. Vol. 1 Amsterdam: Elsevier Science pp. 167–216.
- Luxembourg Income Study (LIS) Microdatabase. 2003. Harmonization of original series conducted by the Luxembourg Income Study, asbl. Luxembourg, periodic updating.
- Lynch, John, George Davey Smith, Marianne Hillemeier, Mary Shaw, Triveliore Raghunathan and George Kaplan. 2001. "Income inequality, the psychosocial environment, and health: comparisons of wealthy nations." *The Lancet* 358:194–200.
- Lynch, John, George Davey Smith, Sam Harper, Marianne Hillemeier, Nancy Ross, George A. Kaplan and Michael Wolfson. 2004. "Is Income Inequality a Determinant of Population Health? Part 1. A Systematic Review." *The Milbank Quarterly* 82(1):5–99.
- Maddison, Angus. 2001. *The World Economy: A Millennial Perspective*. Paris: OECD Development Centre.
- Modigliani, Franco and Richard Brumberg. 1954. Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data. In *Post-Keynesian Economics*, ed. Kenneth K. Kurihara. New Brunswick, NJ: Rutgers University Press pp. 388–436.


- Murray, Christopher J. L., Julio Frenk and Emmanuela E. Gakidou. 2001. Measuring Health Inequality: Challenges and New Directions. In *Poverty, Inequality and Health: An International Perspective*, ed. David A. Leon and Gill Walt. New York: Oxford University Press pp. 194–216.
- Oeppen, Jim and James W. Vaupel. 2002. “Broken Limits to Life Expectancy.” *Science* 296:1029–1031.
- Pampel, Fred C. 2002. “Cigarette Use and the Narrowing Sex Differential in Mortality.” *Population and Development Review* 28(1):77–104.
- Preston, Samuel H. and Paul Taubman. 1994. Socioeconomic Differences in Adult Mortality and Health Status. In *The Demography of Aging*, ed. Linda G. Martin and Samuel H. Preston. Washington: National Academy Press pp. 279–318.
- Rimm, Eric B. and Meir J. Stampfer. 2004. “Diet, Lifestyle, and Longevity — The Next Steps?” *Journal of the American Medical Association* 292(12):1490–1492.
- Roberts, Stephen J. and Will D. Penny. 2002. “Variational Bayes for Generalized Autoregressive Models.” *IEEE Transactions on Signal Processing* 50(9):2245–2257.
- Rogot, Eugene, Paul D. Sorlie, Norman J. Johnson and Catherine Schmitt. 1992. “A Mortality Study of 1.3 Million Persons by Demographic, Social, and Economic Factors: 1979-1985 Follow-Up.” Bethesda, Md: National Institutes of Health; NIH publication 92-3297.
- Rogot, Eugene, Paul D. Sorlie, Norman J. Johnson, Claudia S. Glover and David Treasure. 1988. “A Mortality Study of One Million Persons by Demographic, Social, and Economic Factors: 1979-1981 Follow-Up.” Bethesda, Md: National Institutes of Health; NIH publication 88-2896.
- Shkolnikov, Vladimir M., Evgueni E. Andreev and Alexander Z. Begun. 2003. “Gini coefficient as a life table function: computation from discrete data, decomposition of differences and empirical examples.” *Demographic Research* 8(11).

- 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.
- Smith, James P. 2003. "Consequences and Predictors of New Health Events." *NBER Working Paper* 10063.
- Thomas, Vinod, Yan Wang and Xibo Fan. 2000. "Measuring Education Inequality: Gini Coefficients of Education." *World Bank Working Paper* 2525.
- Tuljapurkar, Shripad and Carl Boe. 1998. "Mortality Change and Forecasting: How Much and How Little Do We Know?" *North American Actuarial Journal* 2(4):13–47.
- Waldron, Ingrid. 1995. Contributions of biological and behavioural factors to changing sex differences in ischaemic heart disease mortality. In *Adult mortality in developed countries: from description to explanation*, ed. Alan D. Lopez, Graziella Caselli and Tapani Valkonen. Oxford: Clarendon Press pp. 161–178.
- Waldron, Ingrid. 2000. Trends in gender differences in mortality: Relationships to changing gender differences in behaviour and other causal factors. In *Gender Inequalities and Health*, ed. Ellen Annandale and Kate Hunt. Buckingham, UK: Open University Press pp. 150–181.
- White, Kevin M. 2002. "Longevity Advances in High-Income Countries." *Population and Development Review* 28(1):59–76.
- Wilkinson, Richard G. 1992. "Income distribution and life expectancy." *British Medical Journal* 304(6820):165–169.
- Wilmoth, John R. and Aaron Gullickson. 2001. "American Mortality in International Perspective: The Evolution of a Distinctive Pattern." Unpublished manuscript, University of California at Berkeley.
- Wilmoth, John R. and Shiro Horiuchi. 1999. "Rectangularization Revisited: Variability of Age at Death within Human Populations." *Demography* 36(4):475–495.

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


World Health Organization Mortality Database. 2004. Available at www.who.int/whosis/mort.

Figure 1: Distributions of ages at death in Sweden and the U.S., 1999



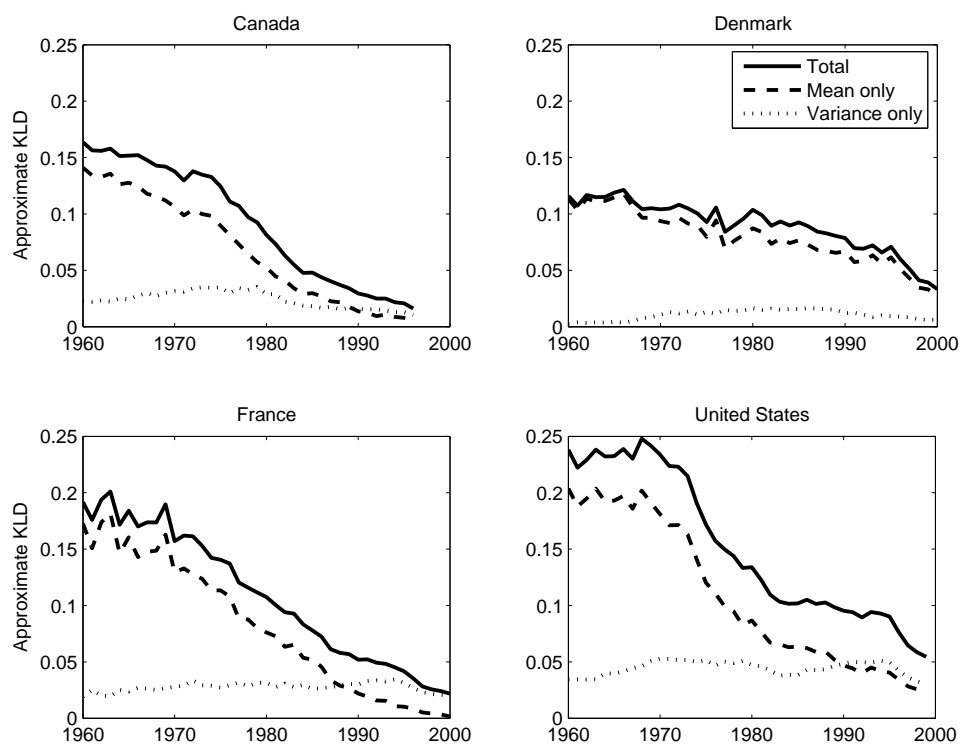
Notes: Data are the ${}_n d_x$ densities from life tables for both sexes combined provided by the Human Mortality Database.

Figure 2: Convergence in 7 adult age-at-death distributions since 1960




Notes: Each line shows the Kullback-Leibler divergence (KLD) for a particular country between its age-at-death distribution for ages 10 and over in a particular year and that of Sweden in 2002. The KLD is described in note 2. Life tables are provided by the Human Mortality Database; age is top-coded at 100.

Figure 3: Convergence in adult age-at-death distributions decomposed into contributions from means and variances




Notes: Data are provided by the Human Mortality Database. Each solid line shows an approximate Kullback-Leibler divergence (KLD) for the country shown in the panel between its age-at-death distribution above age 10 over time and that of Sweden in 2002. The KLD approximation is based only on means and variances and is described in note 3. Each dashed line shows the component of the KLD approximation attributable to differences in the means, i.e. when the variances are equal. Each dotted line shows the component attributable to differences in the variances.

Figure 4: Unconditional standard deviations in the age at death, S_0 , among 7 high-income countries since 1960




Notes: Data are the square roots of variances of ages at death. The weights are life table deaths, ${}_n d_x$, for both sexes combined from the Human Mortality Database.

Figure 5: Conditional standard deviations in the age at death, S_{10} , among 7 high-income countries since 1960




Notes: Data are the standard deviation in ages at death above age 10 for both sexes combined from the Human Mortality Database. The weights are life table deaths, ${}_n d_x$.

Figure 6: Record-low S_{10} among countries in the Human Mortality Database




Notes: Data are the lowest standard deviations of ages at death above age 10 for both sexes combined among the set of 14 high-income countries in the Human Mortality Database in each year depicted. The weights are life table deaths, ${}_n d_x$.

Figure 7: Standard deviations in the age at death conditional on reaching age 10, S_{10} , in the U.S. by race, in Canada, and in France




Notes: Data are standard deviations of ages at death above age 10 for both sexes combined. The weights are life table deaths, ${}_n d_x$, from the Human Mortality Database (France, Canada) and the Berkeley Mortality Database (U.S.).

Figure 8: S_{10} by income group in the U.S. around 1981




Notes: Data are standard deviations of the age at death above age 10 constructed from a life table derived from deaths observed in the first year of the U.S. National Longitudinal Mortality Study. Income was observed at the beginning of the period. M_{10} is the mean age at death above age 10, equal to $e_{10} + 10$. Data have been smoothed using a kernel density estimator.

Figure 9: S_{10} by educational group in the U.S. around 1981



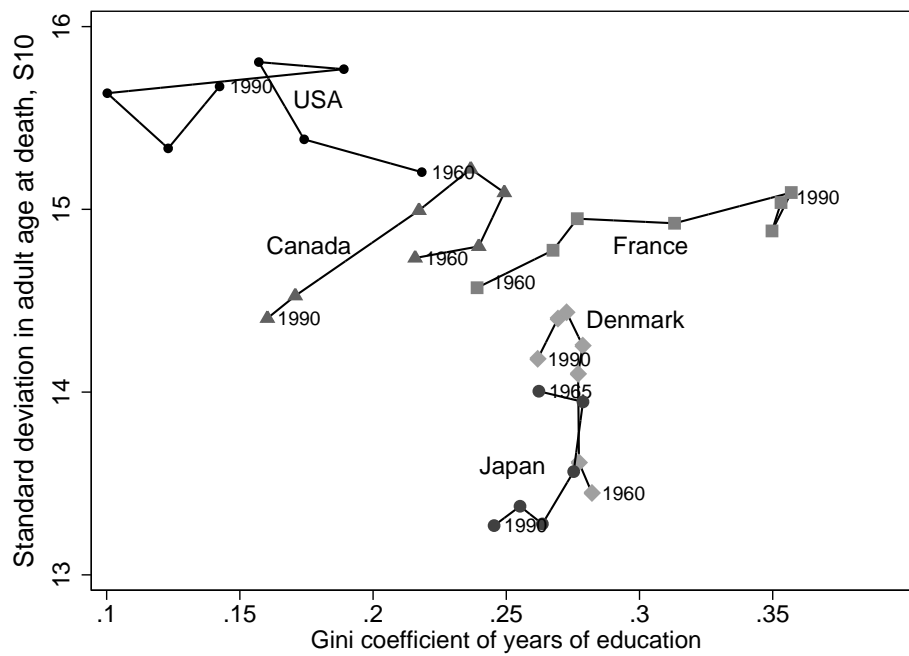
Notes: Data are standard deviations of the age at death above age 10 constructed from a life table derived from deaths observed in the first year of the U.S. National Longitudinal Mortality Study. M_{10} is the mean age at death above age 10, equal to $e_{10} + 10$. Data have been smoothed using a kernel density estimator.

Figure 10: Income inequality and variability in adult ages at death



Notes: The vertical axis measures standard deviations in ages at death above age 10, S_{10} , based on data from the Human Mortality Database. The horizontal axis measures the Gini coefficient of income, provided by the Luxembourg Income Study (LIS) Micro-database (2003), supplemented with Japanese data from Deininger and Squire (1996).

Figure 11: Educational inequality and variability in adult ages at death



Notes: The vertical axis measures standard deviations in ages at death above age 10, S_{10} , based on data from the Human Mortality Database. The horizontal axis measures the Gini coefficient of education, provided by Thomas, Wang and Fan (2000).