

This PDF is a selection from a published volume from the National Bureau of Economic Research

Volume Title: Demography and the Economy

Volume Author/Editor: John B. Shoven, editor

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-75472-3

ISBN13: 978-0-226-75472-7

Volume URL: <http://www.nber.org/books/shov08-1>

Conference Date: April 11-12, 2008

Publication Date: November 2010

Chapter Title: The Final Inequality: Variance in Age at Death

Chapter Authors: Shripad Tuljapurkar

Chapter URL: <http://www.nber.org/chapters/c8414>

Chapter pages in book: (209 - 221)

The Final Inequality

Variance in Age at Death

Shripad Tuljapurkar

6.1 Introduction

Demography and economics shape many aspects of the lives and decisions of individuals as well as the structure and welfare of populations. An important and persistent demographic shift that occupies much attention around the world is the aging of many national populations, driven by changes in the rates of birth, death, or migration. An ongoing decline in death rates is a common factor that drives aging in all industrialized nations and many of the world's developing regions. Birth rates and migration also influence aging, but their importance varies between countries. The twentieth century was the first period in history in which humans experienced a sustained decline in death rates that resulted, in the now-rich nations, in a doubling of human life expectancy at birth and a 50 percent increase in the remaining life expectancy of people at age sixty-five. These changes expanded human life cycles in time and precipitated changes in the pattern of individual lives and in relationships between generations. Economic and demographic analyses of aging work at one or both of these levels. For individuals and families, the stretching of lives affects decisions about the level and timing of life-cycle events such as schooling, work, savings, and retirement. For populations, aging has meant changes in flows of labor and money and challenges related to education, annuities and pensions, insurance, and health care. Analyses at both levels require an understanding of how long people live, the differences between individuals in life spans, and

Shripad Tuljapurkar is the Morrison Professor of Population Studies and a professor of biological sciences at Stanford University.

This research has been supported by grants from the U.S. National Institute of Aging and by the Morrison Institute for Population and Resource Studies at Stanford University.

the rates at which these are changing. One dimension of mortality that has been extensively studied is life expectancy, the average span of life, which is the key statistic used to describe mortality and health conditions. Many studies have examined trends and forecasts of life expectancy, while others have examined the effect of inequalities in wealth, income, or education on health by studying differences in life expectancy between groups that differ in these characteristics.

This chapter focuses on a second dimension of mortality, the variation in lifespan between individuals and groups of individuals. We begin by asking whether the length of life should be measured starting at birth or at some later age. To answer this question, we first show that in today's industrialized countries, childhood mortality is so low that we should focus on differences in the length of adult life. To measure such differences, we define the age at adult death and its variance, following Edwards and Tuljapurkar (2005). This variance and aggregate life expectancy describe two distinct dimensions of the distribution of life (and death) within populations. Next, we present and discuss historical trends in this variance and compare trends across countries. We then discuss the relationship between the pattern of adult death and socioeconomic inequalities, in factors such as education and income, using data from the United States. Finally, we examine the effect of variance in adult death on simple economic measures in an overlapping generations setting.

6.2 Death and Inequality

The modern rise in the length of life began about the time of the Industrial Revolution and has continued ever since. Figure 6.1 illustrates the gains in life expectancy at birth (e_0) and at age sixty-five (e_{65}) using data for Sweden from 1950 to 2000. Over that period, e_0 increased by about 12 percent and e_{65} by about 33 percent. Mortality here is measured using period death rates observed in particular calendar years; for each year, we compute quantities such as the average age at death that describe a hypothetical cohort of individuals who experience those over their lives. The higher proportional increase in e_{65} compared to e_0 resulted from two factors. First, mortality in Sweden at young ages is now so low that further reductions have relatively little leverage on life expectancy. Second, reductions in mortality are, over time, occurring at older ages than in the past. To gain further insight into these two factors, we next examine the probability distribution of the age at death.

The age pattern of mortality is described by an age-specific mortality rate $\mu(a)$, and the probability of living to at least age a is the survivorship $l(a)$. The probability that an individual dies at age a is described by the density $\phi(a) = \mu(a)l(a)$. Figure 6.2 displays this density for Sweden in 1950 and in 2000. The risk of dying at young ages is concentrated in the first year of life




Fig. 6.1 Gains in period life expectancy between 1950 and 2000 at birth (e_0 , solid) and at age 65 (e_{65} , dashes) for Sweden, both sexes combined



Fig. 6.2 Probability distribution of age at death in 1950 (solid) and 2000 (dashed) for Sweden, both sexes combined

and has fallen steadily in the past fifty years. For example, in Sweden in 2000, less than 0.4 percent of deaths in the period life table occur at ages under ten years. Beyond age ten, death is increasingly likely, with over 85 percent of all deaths concentrated in a range of twenty years or so around a sharply defined modal age that is slightly higher than the life expectancy at birth. It is the variation in this age range that describes the bulk of variation in “adult” death. An individual who survives his or her first year of life is most likely to die as an adult (over age ten), and differences between individual ages at death are largely differences in the age of adult death.

Based on these observations, Edwards and Tuljapurkar (2005) define adult death as death occurring after age ten. The probability distribution of the age of adult death is derived from $\phi(a)$ in figure 6.2 as the conditional distribution given that death occurs after age ten. The shape of the conditional distribution is the same as that of ϕ . The variance of this conditional distribution is defined to be the variance in the age at adult death, denoted here by S_{10}^2 . The value of S_{10} measures the dispersion in age at adult death. We cannot measure this dispersion by the variance of the full distribution ϕ because the size of that variance is always strongly affected by the infant mortality peak even when infant mortality is as small as it is in figure 6.2. Our choice of ten years is somewhat arbitrary, but any age near the minimum of the full distribution (see figure 6.2) serves equally well. Figure 6.3 shows the effect of using different cutoff ages of ten and twenty years on the standard deviation of the age at adult death, using data for Sweden from 1951 to 2000. The two curves shown track each other closely, and the values are very close over the period.

The measure S_{10} describes the extent of inequality in the age at death. Why do we call this an inequality? There is considerable current interest in the role of socioeconomic inequalities as determinants of inequalities in health outcomes (e.g., Marmot 2005). Health is not easily defined or measured, but mortality risk is widely used as an indicator of health, and age at death is, of course, a primary health outcome variable. In this context, our S_{10} is an appropriate measure of inequality in health outcomes. We note that a different way of describing inequality in adult death is to use percentiles of the death distribution, as suggested by Victor Fuchs in his comments on this chapter. Such percentiles have previously been used by Wilmoth and Horiuchi (1999) in a discussion of the possible compression of age at death. We believe that S_{10} is in many ways a natural measure and is particularly useful in thinking about the nature of risk, but percentiles can provide useful additional insights.

The distribution of adult deaths is the large concentrated mass of the distribution in figure 6.2. A rough approximation to the distribution is a normal centered on the modal age at death with a standard deviation of S_{10} , and we use this approximation later in this chapter. It is worth comparing the actual distributions in figure 6.2, or their normal approximations, to two stylized distributions of death that have been used by economists. The first, dating

back to early work (Yaari 1965; Blanchard 1985) on overlapping generation models, assumes that the probability of death is independent of age (panel A of figure 6.4) and leads to a most unrealistic exponential distribution of the age at death. The second (Futagami and Nakajima 2001) assumes that all adults die at the same age (panel B of figure 6.4). Our discussion suggests that a more realistic treatment of the age distribution of human deaths




Fig. 6.3 The effect of defining “adult” death as deaths over age 10 or 20. The solid line shows S_{10} , and the dashed line shows S_{20} , as defined in the text, for Sweden from 1950 to 2000, both sexes combined




Fig. 6.4 Stylized probability distributions of age at death: *A*, Age-independent probability of death; *B*, All deaths at one age

should use e_0 , which is close to the modal age of adult death, as a measure of location and S_{10} as a measure of dispersion.

6.3 Historical Inequality in Adult Death

Historical changes have increased the average age at death e_0 in most countries. We now examine the corresponding historical change in the dispersion in adult death measured by S_{10} . The nature of change in S_{10} will tell us whether mortality improvement means that both the average and the variance in adult age at death change together. In other words, are we compressing inequality in age at adult death while also delaying death?

Figure 6.5 plots S_{10} versus life expectancy e_0 for Sweden from 1900 to 1950. Time turns out to run from left to right across the plot. There were fluctuations in both e_0 and S_{10} , but the overall negative correlation between them was very high. In this period, S_{10} fell to 50 percent of its 1951 value, decreasing at 0.22 years per calendar year, whereas e_0 grew to nearly 150 percent of its value in 1951, increasing at 0.4 years per calendar year. In the years 1951 to 2000, as shown in figure 6.6, the negative correlation between S_{10} and e_0 weakened somewhat. Life expectancy continued to increase, albeit at a slower pace, at about 0.2 years per calendar year. But S_{10} decreased much more slowly and with significant fluctuation, at about 0.022 years per calendar year.

In the first half of the twentieth century, mortality declines clearly acted as a “rising tide” that reduced inequality in age at adult death across the




Fig. 6.5 Standard deviation S_{10} in adult age at death plotted against life expectancy at birth e_0 from 1900 to 1950 for Sweden, both sexes combined




Fig. 6.6 Standard deviation S_{10} in adult age at death plotted against life expectancy at birth e_0 from 1951 to 2000 for Sweden, both sexes combined

population as a whole. In terms of the distribution of age at death (recall figure 6.2), the mass of adult deaths moved to later ages while also being compressed. In the second half of the twentieth century, progress against mortality continued, so the mass of deaths continued its march to older ages, but the compression of inequality slowed considerably. It is important to recognize that the compression of mortality inequality contains an important message about the extent of variation in mortality between individuals. There is great interest in the effect of risk factors as predictors of individual mortality risk, and the notion that individual behavior can strongly affect age at death is widespread. Indeed, the argument is often made that the distribution of risk factors shapes the distribution of deaths (e.g., Mokdad et al. 2004). History tells us, however, that the total variance in adult death, which includes the contributions of all risk factors, has declined substantially over time and indeed continues to do so. We return to the predictive value of risk factors later in this chapter.

6.3.1 International Trends and the Future

How do these historical patterns for Sweden compare with what has happened in other countries? The slowdown in the decline of S_{10} in Sweden since about 1960, seen in figures 6.5 and 6.6, is partially mirrored across the industrialized world. A comprehensive and recent comparison across all Organization for Economic Cooperation and Development (OECD) countries has been published by the OECD (2007). We focus on a subset of the OECD countries from 1960 onward as shown in figure 6.7, which is redrawn




Fig. 6.7 Conditional standard deviations in the age at death, S_{10} , in seven high-income countries since 1960

Source: Data for both sexes combined are taken from the Human Mortality Database (<http://www.mortality.org>).

from the data used by Edwards and Tuljapurkar (2005). The strikingly highest and steadiest curve in the plot is for the United States, which had the highest level of mortality inequality among these countries (and, indeed, across the industrialized world) over the entire period. Canada displayed a level of inequality and a lack of trend similar to the United States from 1960 to 1980, but after that, S_{10} in Canada has fallen significantly. The sharp contrast between recent trends in S_{10} in these two countries is plausibly due to the widespread availability of national health services in Canada after 1980. For the entire period shown in figure 6.7, there is one country whose S_{10} is just below that for the United States and shows the same absence of overall trend. That country is France. Given the widespread public commentary in each country that they are least likely to resemble each other, this is quite a surprise.

The United Kingdom, Sweden, and Denmark started out with similar levels of inequality in 1960. Sweden and the United Kingdom changed little through the 1980s, but Sweden's S_{10} then declined, whereas the United Kingdom had a modest increase. Denmark is another surprise, with an increase

in S_{10} through the 1980s and higher inequality at the end of the period than it had in 1960. Japan, as is often the case in such comparisons, is strikingly distinctive, with a notable decrease in inequality from 1960 (when Japan and the United States had similar levels of S_{10}) till 1990 (when Japan and Sweden were tied with the lowest inequality). In the most recent decade, Japan's S_{10} has actually increased. Victor Fuchs (in his comments on this chapter) has examined this recent trend in Japan using percentiles of the distribution of age at death. To see why percentiles matter, look again at figure 6.2. The distribution of age at death around the mode has a left skew, as is typical of most human history, which means that much of the inequality we discuss here is driven by early deaths. But for recent years in Japan, Fuchs finds that the probability of dying at ages above the mode (use figure 6.2 as a guide) has increased relative to the past, thus changing the skewness of the distribution. As a result the inequality in age at death in Japan may be increasing because there is a higher chance of living to old ages past the mode. This explanation marches with the known fact that the number of centenarians in Japan is increasing very rapidly with time (Robine, Saito, and Jagger 2003).

Bongaarts (2007) recently proposed an interesting model of mortality change to be used in making forecasts. He argues that life expectancy simply increases at some steady rate per year and that the shape of the distribution of adult deaths, based on $\phi(a)$, does not change with time for deaths over age twenty-five years. In his view, the mass of adult deaths, as shown in our figure 6.2, simply translates to later ages at some steady rate, but with the dispersion of the mass constant. He arrived at his model using rather different arguments about the nature of senescence, and so our historical analysis provides a test of his assumptions. It is clear from figure 6.7 that his approximation is plausible for trends in the United States since 1960; it may also be plausible for some other, but not all, countries in recent decades. His model would clearly not be correct as a description of historical change prior to 1960.

6.4 The Sources of Variance in Adult Death

We turn now to a different question: what causes differences in mortality within a country between groups that are distinguished by characteristics such as income, education, race, or other factors that we expect to influence mortality risk? This question has become particularly important in recent discussions about the relationships between mortality and socioeconomic inequality measured in various ways (Mokdad et al. 2004; Marmot 2005). Typically, analyses of such relationships have focused on the effect of a particular risk factor on either life expectancy or relative mortality rates. Controlling for differences in other likely risk factors, a successful analysis detects a difference in the e_0 corresponding to differences in the particular factor in question. Such studies measure what we call the variance *between*

groups that are distinguished by particular explanatory factors. But we have found that such relationships can be studied in a different and more informative way by asking how socioeconomic factors affect the variance of adult age at death both *between* groups and *within* groups.

We consider a decomposition of a population into subgroups based on differences in socioeconomic variables and use results from Edwards and Tuljapurkar (2005). They considered the effects of education and income, both factors that are well known to be correlated with mortality rates and average age at death, as well as of sex, race, and certain causes of death. We focus on the effects of education, which is a much more stable socioeconomic measure for adults than is income. Data were taken from the U.S. National Longitudinal Mortality Study, a panel study of over half a million individuals who were interviewed around 1980 and then tracked for nine years. Socioeconomic data were observed only at the beginning of the period, and the analysis used only mortality in the first year of the sample. To keep comparisons simple, the analysis considered only two socioeconomic strata, with individuals sorted according to whether they are high school graduates, roughly two-thirds of the sample. Life tables were constructed for both sexes combined in each group, and smoothed distributions of ages at death were constructed and used to estimate conditional means and variances.

Figure 6.8 (redrawn using the data from Edwards and Tuljapurkar 2005) plots distributions of age at adult death by educational status. The plot lists for each group the values of the conditional mean age at death M_{10} and the within-group standard deviation S_{10} . Clearly, adults in the lower stratum not only have shorter average life spans, but also are subject to greater variability. As adults, high school graduates live an average of five years longer than their less-educated counterparts, while enjoying a standard deviation that is two years lower. But the variance between these groups (approximately the square of the difference in M_{10} , so ≈ 25) is an order of magnitude smaller than the variance within groups (the average of the variances, so ≈ 225). This huge difference reflects the considerable overlap between the two distributions in figure 6.8. Even if everyone in the United States had a high school diploma, S_{10} would remain fairly high, at 14.6, which is only a year lower than the value for the United States as a whole. Clearly, education matters, but it matters more to averages and rather less to inequality, and, thus, matters less to the predictive power of education about the age of death. A similar result is found when looking at age at death as a function of household income (Edwards and Tuljapurkar 2005).

These results lead to broad conclusions about analytical strategies for future research and about policy conclusions from existing research. The analytical strategy used to study the effects of socioeconomic inequality needs to focus on mortality inequality and not just on average outcomes. For example, it would be useful to search for risk factors that best separate groups, that is, that maximize the ratio of between-group variance to within-




Fig. 6.8 Distributions of ages at death by educational group in the United States in 1981

Source: Data are constructed from a life table derived from deaths observed in the first year of the U.S. National Longitudinal Mortality Study.

Notes: Education was observed at the beginning of the period. M_{10} is the mean age at death above age ten, equal to $e_{10} + 10$. Data have been smoothed using a kernel density estimator.

group variance in adult age at death. It would be useful to ask whether the roughly constant inequality in age at death in the United States can be explained by changes in socioeconomic inequality. In other countries where S_{10} has fallen over time, we should ask whether the effect of mortality decline has been to reduce the within-group variances for all groups or just the variances within particular groups. In terms of policy, the results show clearly that reducing some kinds of socioeconomic inequality will have little or no effect on inequality in age at death.

6.4.1 Economic Theory and Variance in Adult Death

Our variance S_{10} is simply the dispersion of the random age at death, call it T , across adult individuals in a population. We can approximate the distribution of adult deaths by a normal distribution around the modal age at death, call it μ , with a standard deviation $\sigma = S_{10}$. This approximation undershoots the true left-skewed distribution at ages below μ and overshoots the true

distribution at ages much over μ , but it is reasonable for seeing how variance in T affects lifetime income, consumption, and utility.

Suppose that wages are fixed at some value W and an individual works starting at some age a_s (upon leaving school or college, say) until the earlier of death or retirement at age a_r . For a given interest rate r , expected lifetime earnings are

$$I = WE \int_{a_s}^{(T \wedge a_r)} ds e^{-rs} = \left(\frac{W}{r} \right) [e^{-ra_s} - Ee^{-r(T \wedge a_r)}].$$

Here E indicates an expectation over the distribution of age at death T , which we take to be a normal distribution as in the preceding. The exact expressions here are messy, but they are closely approximated by

$$I = \left(\frac{W}{r} \right) \{ e^{-ra_s} - l(a_r) e^{-ra_r} - [1 - l(a_r)] e^{-r\mu + (1/2)r^2\sigma^2} \}.$$

This is sensible: when retirement occurs at an age well below the modal age at death μ , uncertainty in death has little effect on lifetime income. As age at retirement increases toward μ , the dispersion σ in T translates into dispersion in lifetime income. There is a trade-off between μ and σ , in that

$$\frac{\partial I}{\partial \sigma} = -r\sigma \frac{\partial I}{\partial \mu}.$$

For an interest rate of 0.03, and σ ; 14, which is typical of industrialized countries, the multiplier is 0.42; in developing countries with σ ; 25, the multiplier is 1. So the effect of increasing μ by a year is about the same as decreasing σ by half a year in industrialized countries and by a year in developing countries.

Lifetime consumption also depends on T . In simple overlapping generations models (Blanchard 1985) with constant relative risk aversion (CRRA) utility, the optimal consumption at age x is a function

$$c(x) = c_0 e^{kx}, \text{ where } k = \frac{(r - \theta)}{\gamma},$$

where r is interest rate, θ is the discount rate, and γ is the coefficient of risk aversion. Lifetime consumption then depends on e^{kT} , and we have

$$Ee^{kT} = e^{k\mu + (1/2)k^2\sigma^2}.$$

So inequality in T translates into inequality in lifetime consumption. This fact suggests that it would be useful to incorporate uncertainty in T into analyses of the benefits of increasing lifespan.

Lifetime utility depends on consumption in these settings, and in the CRRA model, utility at age x is proportional to $c(x)^{(1-\gamma)/(1-\gamma)}$. Expected lifetime utility averages over the variation in T and, thus, also depends on

σ . The effect of σ on lifetime consumption depends on the factor k , but the effect on lifetime utility depends on the product $k(1 - \gamma)$, being modified by the level of risk aversion. Li (2005) has explored these connections in more detail by studying the equilibrium of a simple closed economy model with adult deaths distributed normally as in the preceding.

6.5 Conclusion

This chapter has shown that the variance in age at adult death is a useful and important dimension of mortality change. Trends in this variance are informative about the speed and the age pattern of mortality change. The decomposition of this variance with respect to risk factors provides useful insights into the explanatory power of different factors that are correlated with mortality. Historical and economic analyses can benefit from an examination of variance in age at death in addition to the traditionally important study of life expectancy.

References

- Blanchard, O. J. 1985. Debt, deficits and finite horizons. *Journal of Political Economy* 93:223–47.
- Bongaarts, J. 2005. Long-range trends in adult mortality: Models and projection methods. *Demography* 42:23–49.
- Edwards, R., and S. Tuljapurkar. 2005. Inequality in life spans and a new perspective on mortality convergence across industrialized countries. *Population and Development Review* 31 (4): 645–74.
- Futagami, K., and T. Nakajima. 2001. Population aging and economic growth. *Journal of Macroeconomics* 23:31–44.
- Li, Q. 2005. Essays on asset pricing, consumption and wealth. PhD diss., Stanford Univ.
- Marmot, M. 2005. Social determinants of health inequalities. *Lancet* 365:1099–1104.
- Mokdad, A. H., J. S. Marks, D. S. Stroup, and J. L. Gerberding. 2004. Actual causes of death in the United States, 2000. *Journal of the American Medical Association* 291:1238–45.
- OECD. 2007. *Society at a glance: OECD social indicators—2006 edition*. Paris: Organization for Economic Cooperation and Development.
- Robine, J.-M., Y. Saito, and C. Jagger. 2003. The emergence of extremely old people: The case of Japan. *Experimental Gerontology* 38:735–39.
- Wilmoth, J. R., and S. Horiuchi. 1999. Rectangularization revisited: Variability of age at death within human populations. *Demography* 36:475–95.
- Yaari, Menahem E. 1965. Uncertain lifetime, life insurance, and the theory of the consumer. *Review of Economic Studies* 22:137–50.