NBER WORKING PAPER SERIES

VARIANCE IN DEATH AND ITS IMPLICATIONS FOR MODELING AND FORECASTING MORTALITY

Shripad Tuljapurkar
Ryan D. Edwards

Working Paper 15288
http://www.nber.org/papers/w15288

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2009

Tuljapurkar: Morrison Professor of Population Studies, Department of Biological Sciences, Stanford University, tulja@stanford.edu. Edwards: Assistant Professor of Economics, Queens College and the Graduate Center, City University of New York, and the National Bureau of Economic Research, redwards@qc.cuny.edu.
Tuljapurkar acknowledges support from NIA grant 1P01 Edwards acknowledges support from the Morrison Institute and from NIA grant T32 AG000244. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

Variance in Death and Its Implications for Modeling and Forecasting Mortality
Shripad Tuljapurkar and Ryan D. Edwards
NBER Working Paper No. 15288
August 2009
JEL No. I1,J11,N3

ABSTRACT

Entropy, or the gradual decline through age in the survivorship function, reflects the considerable amount of variance in length of life found in any human population. Part is due to the well-known variation in life expectancy between groups: large differences according to race, sex, socioeconomic status, or other covariates. But within-group variance is very large even in narrowly defined groups, and it varies strongly and inversely with the group average length of life. We show that variance in length of life is inversely related to the Gompertz slope of log mortality through age, and we reveal its relationship to variance in a multiplicative frailty index. Our findings bear a variety of implications for modeling and forecasting mortality. In particular, we examine how the assumption of proportional hazards fails to account adequately for differences in subgroup variance, and we discuss how several common forecasting models treat the variance along the temporal dimension.

Shripad Tuljapurkar
Stanford University
Center on Longevity
Herrin Labs, Room 454
Stanford, CA 94305
tulja@stanford.edu

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

Introduction

Length of life is a fundamental dimension of human prosperity. We measure this dimension either with period life expectancy at birth, e_0, the average length of life, or we measure its inverse either with age-specific mortality rates that underpin the life table, or with the log odds of death. Correctly modeling mortality is crucial for inference in both observational and experimental settings, and thus for forecasting. In this paper, we illustrate how patterns in the variance of length of life, whether measured across subgroups at a point in time or in human populations over long periods of time, bear strong implications for how we model mortality and test hypotheses in cross-sectional and longitudinal settings.

In the cross section and over short panels, the Cox (1972) proportional hazards model and the logit or logistic regression model are standard tools in epidemiological studies and in medical research. As they are typically specified, these models assume that subgroups experience proportionally higher or lower hazards relative to a baseline. We show that cross-sectional patterns in the variance in length of life across subgroups, where variance is inversely related to average life expectancy, are not at all adequately captured by proportional hazards. This is because variance in length of life is closely tied to the age slope of mortality; subgroup differences in the variance in length of life are equivalent to subgroup differences in the age slope of mortality. Regardless of the precise nature of baseline mortality, which may be modeled nonparametrically, proportional hazards impose the same age slope and thus the same variance in length of life on all subgroups.

Violations of the proportional hazards assumption have been extensively remarked and explored (Hess, 1995; Lee and Go, 1997; Therneau and Grambsch, 2000), and researchers have suggested various methods to address them. One option is to augment standard models by injecting individual frailty (Vaupel, Manton and Stallard, 1979; Hougaard, 1995). As we discuss, frailty models are helpful because they add an additional parameter that directly impacts variance. But the concept of frailty does not help us intuitively understand cross-sectional or intertemporal patterns in variance; it simply improves the fit of models without

increasing our understanding. One class of models designed to account for frailty, accelerated failure time (AFT) models, posits a proportional scaling of the distribution of survival time across subgroups, which inappropriately implies a positive rather than negative relationship between mean and variance. Approaches that relax the assumption of proportional scaling, such as the use of stratified Cox regression, time-dependent covariates, or the nonparametric method of Kaplan and Meier (1958) are preferable.

From an aggregate perspective, the expansion of human life in the past century (Preston, 1975; Caldwell, 1976) and its socioeconomic implications have stimulated efforts to analyze and forecast mortality trends (Tuljapurkar and Boe, 1998), which are guided by insights gained from mortality models. A natural focus of these efforts is the period expectation of life at birth, e_0. Mortality change is commonly summarized in terms of trends in e_0, and mortality models are evaluated on their ability to match historical trends in life expectancy. These uses of e_0 gained considerable support from two recent findings: that e_0 has increased at a nearly constant rate in many industrial countries since 1955 (White, 2002), and that since 1840 annual world record female e_0 has also increased at a nearly constant rate (Oeppen and Vaupel, 2002). Some have argued that such constancy is fundamental in analyzing mortality change (Bongaarts and Feeney, 2002, 2003), and one researcher (Bongaarts, 2005) has extended a simple model (Vaupel, 1986) to forecast mortality change. But e_0 is only the mean of the distribution of ages at death, and we show here that the variance of this distribution provides important additional information. As we reveal, temporal change in the variance in age at adult death is not necessarily captured by simple models, which may inappropriately constrain how we should conceptualize and analyze mortality change.

This paper is organized in four main parts. First, we discuss cross-sectional and temporal patterns in distributions of period life-table ages at death. We recount how historical increases in e_0 in the industrialized countries have been accompanied by equally striking decreases in the variance of the age of adult death (Edwards and Tuljapurkar, 2005). These trends show clearly that mortality decline over time has compressed the variance between

individuals at the same time as it has increased average life expectancy. Cross-sectional patterns reveal the same inverse relationship between subgroup variance and average. While overall variance has declined over time, very large differences in variance between subgroups remain. Less advantaged groups experience both lower average length of life and higher variance.

Second, we show how the variance in age of adult death can be approximately computed for any reasonable model of mortality rates, and illustrate this with three commonly used models, the Gompertz (1825), the logistic, and a Gompertz model with multiplicative frailty. In particular, we reveal the inverse relationship between the Gompertz slope and the variance in length of adult life. A particular age slope of mortality will reflect a particular variance, but while it can be consistent with different subgroup average lengths of life, it cannot accurately capture the large differences in subgroup variances. Adding frailty into the model can explain variance, but not in a particularly satisfactory way in terms of any intuition.

Third, we explore the implications of these insights for common mortality models. As concerns forecasting, we show that any generalization of the Bongaarts-Vaupel translation argument yields an unchanging variance in the age at adult death over time, which may or may not be a preferable characteristic. Over long periods of time, trends in the variance have comprised a major qualitative aspect of mortality change in industrialized countries, and we suspect the same is true for developing countries. We recount world-record trends in the variance in age of adult death, and discuss their implications for understanding secular mortality change. In the setting of panel data on individuals, we illustrate how Gompertz slopes that vary systematically across subgroups violate the assumption of proportionality, and we discuss how various methods to address this either succeed or fail in modeling variance correctly.

Finally, we provide a more general discussion of how our results fit into ongoing research on aging more generally. Insights into the plasticity of the Gompertz slope and what may be driving it are particularly relevant for gauging the future of the human aging process. Results

in the literature on genetic interventions in nonhuman species become particularly intriguing when combined with the insights we present here. Our results are related to work on the "rectangularization" of the survivorship schedule (Wilmoth and Horiuchi, 1999; Kannisto, 2000) and on its shape more generally (Cheung et al., 2005). They are also connected to research on the existence of a maximum age at death (Fries, 1990; Olshansky, Carnes and Cassel, 1990; Wachter and Finch, 1997). All these topics subsume questions as to the possible limiting forms of the distribution of age at death. Our analysis makes no assumptions or deductions about such a limit, but aims to illuminate the nature and significance of trends in the variability of age at death.

Distributions of Age at Death Over Time, Space, and Characteristic

The age pattern of life table deaths in any period, which is also the distribution of period length of life, is found from mortality rates $\mu(a)$ by age a in that period. We define cumulative mortality as

$$M(a) = \int_0^a \mu(s)\ ds, \tag{1}$$

and survivorship is given by

$$\ell(a) = \exp\left[-\int_0^a \mu(s)\, ds\right] = \ell(a) = \exp[-M(a)]. \tag{2}$$

The probability density of death at age a is

$$\phi(a) = \mu(a)\ \ell(a). \tag{3}$$

Even when it is negligible, such as across industrialized countries during the past half century, infant mortality produces some nonzero $\phi(0)$ at the extreme left end of the distribution. But

$\phi(a)$ falls to zero thereafter and remains low; the majority of deaths occur at much later ages.

Figure 1 plots $\phi(a)$ for racial subgroups within the U.S. in 2004 based on life tables prepared by Arias (2007). In that year, period life expectancy at birth for African Americans was 73.1 years, 5.2 fewer years than for whites. Part of this widely-remarked racial inequality was due to a greater density of deaths in infancy, which approached 0.015% for blacks, more than twice the level for whites. But there were also racial differences in the much larger probability of death at older ages. Life expectancy conditional on reaching age 10, e_{10}, was 68.9 for whites but only 64.3 for blacks, a gap that died out gradually above age 25. The width of the distribution around older ages, which we measure by the standard deviation above age 10 or S_{10}, introduced by Edwards and Tuljapurkar (2005) and discussed shortly, is also visibly different across racial groups. For whites in 2004, $S_{10} = 14.9$, while among blacks it was almost 2 years higher. Patterns of inequality in length of life through other dimensions of socioeconomic status (SES) such as income or education look essentially the same (Edwards and Tuljapurkar, 2005), as are differences by sex, although they are more muted.

A very similar picture emerges when we examine distributions of length of life over time instead of across characteristic. Figure 2 plots deaths for both sexes combined in Sweden in 1900 and 2000 using the same scaling of the vertical axis as in Figure 1. These data from the Human Mortality Database (2009), which have been smoothed less than those in the NCHS life tables, reveal that long-term temporal variation in the distribution looks a lot like cross-sectional variation: higher status or more time is rewarded with less variance in addition to a higher mean. Dissimilarities include the much higher level of infant mortality in 1900, which is literally off the chart in Figure 2, and a considerable "baseline" probability of death at practically any age. Both of these characteristics are common to mortality patterns before the completion of the epidemiological transition, when infectious disease makes death a significant probability at any age but especially at birth. This is a large part of the reason

why Swedish S_{10} was 20.7 that year, compared with 12.9 in 2000.

Comparing Figures 1 and 2 reveals another pattern, namely that there are interesting differences in the distribution of length of life across geographic boundaries. The standard deviation in adult length of life among U.S. whites, who are roughly representative of the nation as a whole, is currently higher by about 2 years than it is in Sweden, where tighter concentration pushes the mode higher. This is the point made by Wilmoth and Horiuchi (1999) and Edwards and Tuljapurkar (2005), and recently extended to developing countries by Edwards (2009*c*). Cross-national trends conform to the now familiar pattern of higher status bringing higher average life expectancy and lower variance in the broad cross section of countries over time, but among industrialized countries the relationship is less clear (Edwards and Tuljapurkar, 2005).

We have thus far focused on variance in adult ages at death, choosing a cutoff age A that separates infant and early childhood deaths from later deaths. A suitable value of A lies in the range of ages at which probabilities of death are near their minimum and are relatively stable over time, but is otherwise arbitrary. Following Edwards and Tuljapurkar (2005), we focus on S_{10}, the standard deviation of length of life starting from age 10, but we find results to be similar across various cutoff ages. Still, there is controversy over whether specifying any cutoff age at all is appropriate. Although infant and child deaths contribute strongly to total variation in length of life, we find that they do so in a relatively uninformative way that masks other important trends.

How do young deaths (at ages $\leq A$ years) and adult deaths (at ages $> A$) contribute to mortality? Consider first the average age at death in the period life table, also known as period life expectancy at birth. Write T for the random age of death of an individual in a hypothetical cohort following a period life table. Let p_-, p_+ be the probabilities of young death ($T \leq A$) and adult death ($T > A$) respectively. Then period life expectancy at birth can be decomposed as

$$e_0 = p_- \, M_- + p_+ \, M_+, \tag{4}$$

where M_- and M_+ are (conditional) average ages of death for those who die young or die as adults, respectively. In the industrialized countries in the last five or six decades, M_- is much below 1 year for all subgroups, and p_- is well under 10%, so the main determinant of e_0 is the timing of adult death. Of course, declines in young deaths still matter to e_0, but their effect is proportional to the value of M_+. Consistent with this observation, Wilmoth and Horiuchi (1999) used different methods to show that mortality change at adult ages has been the main contributor to changing e_0 in recent decades.

Now consider the variance in period length of life. We can decompose total variance starting from birth as

$$\mathrm{Var}(T) = p_- \, V_- + p_+ \, V_+ + p_- \, (M_- - e_0)^2 + p_+ \, (M_+ - e_0)^2, \tag{5}$$

where V_- and V_+ are (conditional) variances of age at death for those who die young or die as adults, respectively. For all subgroups in the industrialized countries, only the second and third of these four terms matter. The first term is small because both its components are small, and the last term is small because e_0 has become almost arbitrarily close to M_+. While the third term contributes substantially to $\mathrm{Var}(T)$, it does so only because M_- is very small relative to e_0. This difference is not at all informative about substantive variation in the adult ages at which most deaths occur.

To illustrate, we show in Figure 3 the four components in equation (5) in Swedish data over the period 1900 to 2003, setting our cutoff age A to 10 as usual. Since about 1940, the element that matters most to understanding total variability in age at death is the second term in equation (5), $p_+ \, V_+$, which is shown at upper-right. Today, variance in length of life above age 10 accounts for over 85% of total variability in Sweden, while variance below age 10 such as attributable to infant mortality is responsible for less than 15%. Consistent with this analyis, Edwards and Tuljapurkar (2005) show that differences between industrialized countries in the distribution of age at death are increasingly determined by differences in V_+

rather than in e_0.

While we have not performed a formal analysis, we believe a similar bottom line ought to emerge were we to examine distributions of length of life among subgroups defined by SES or race, which is difficult owing to data constraints. We know it to be true for either sex, and we also find that sex differences in variance today are driven primarily by sex differences in adult variance rather than in infant mortality, a result that mirrors work by Glei and Horiuchi (2007) on the sex differential in life expectancy. Given that trends in variance are important and interesting, a vital question is whether we are capturing them correctly when we model mortality.

Theoretical Models of Adult Mortality

How is the variance in age at adult death described by mortality models? The most celebrated and influential model of adult age-specific mortality is that of Gompertz (1825), in which the force of mortality rises exponentially with age. But recent work by Vaupel et al. (1998), Thatcher, Kannisto and Vaupel (1998) and others suggests that a logistic model with an asymptote describes old-age mortality more accurately. The logistic can also be seen as a result of a model in which Gompertz mortality is modified by a multiplicative frailty (Vaupel, Manton and Stallard, 1979). Frailty, if it occurs in this form, should clearly contribute to the variability in age at death. While it can, we do not find this channel to be an entirely compelling account of historical or cross-sectional patterns in S_{10}. Overall, we will show that traditional models are not well-equipped to deal with or provide understanding about variance, and we find that disturbing in light of its clear importance.

We now present analytical results showing how the variance in age at adult death depends on the parameters of mortality models. We consider in order a general mortality model, the Gompertz, the logistic, a general model with multiplicative frailty, and the Gompertz with multiplicative frailty. We close this section with a summary of our findings

before proceeding in the following section to examine the implications of trends in variance for specific mortality models in use today.

General Mortality Model

Suppose that adult mortality $\mu(a)$ is an increasing positive function of age a. Survivorship falls to zero as age a increases because cumulative mortality $M(a)$ is increasing. The probability distribution of age at death for adults, $\phi(a)$, increases at young adult ages and falls to zero at very high ages. Using an apostrophe to indicate a derivative with respect to age, we differentiate equation (3) to reveal

$$\frac{d\phi}{da} = \phi' = \mu'\,\ell + \mu\,\ell'.$$

The change in the value of ϕ between age a_1 and a slightly larger age $a_1 + x$ is the sum of two terms. The first, $x\,\mu'(a)\,\ell(a)$, represents an increase in the probability due to the increase with age of the death rate μ; the second, $x\,\mu(a)\,\ell'(a) = -x\,\mu^2(a)$, represents a decrease in the probability due to the decrease with age of survivorship $\ell(a)$. At the modal age at death, which we will denote a_0, density is maximized. Because $\phi' = 0$, the two components must be perfectly balanced:

$$\mu'(a_0) = \frac{d\mu}{da} = \mu^2(a_0). \tag{6}$$

If the mortality curve $\mu(a)$ were to steepen so that μ' were higher at every age a, the mode would have to shift to a younger age in order to maintain equality.

Near the mode a_0, the age-at-death distribution $\phi(a)$ can be approximated via Taylor expansion by a quadratic function,

$$\phi(a) \approx \phi(a_0)\left(1 - \frac{(a-a_0)^2}{2\,\sigma^2}\right), \tag{7}$$

where

$$\sigma^2 = \frac{\phi(a_0)}{|\phi''(a_0)|} = \frac{\mu(a_0)}{|\mu''(a_0) - 2\,\mu^3(a_0)|}. \tag{8}$$

Here $\phi''(a_0)$ is the (negative) second derivative of $\phi(a)$ evaluated at the mode a_0 and $\mu''(a_0)$ is second derivative of $\mu(a)$ at the mode a_0. When the distribution $\phi(a)$ is reasonably sharply peaked around the mode a_0, we can approximate it by a normal distribution,

$$\phi(a) \;\approx\; \phi(a_0)\;\exp\left(-\frac{(a-a_0)^2}{2\,\sigma^2}\right). \tag{9}$$

This approximation provides a useful and often accurate estimate of the moments of $\phi(a)$ – we use it here and also check its accuracy by numerical computation (***REF*** asymptotic expansions). In particular, the variance in age at adult death is approximately given by the σ^2 appearing in equation (8). This variance depends on the curvature of the mortality function, i.e., whether the slope of mortality steepens or shallows around the modal age. If the curve steepens, then $\mu''(a_0) > 0$ and the variance is smaller than for a curve that shallows at the mode.

The Gompertz Model

We write the Gompertz mortality function as

$$\mu(a) = \mu_0\, e^{\beta\, a}, \tag{10}$$

where the parameter β is the familiar age-slope of log mortality, a constant in the Gompertz model. In the U.S. for both sexes combined, the age-slope is about 0.087 (Edwards, 2009*a*); that is, mortality rates increase 8.7% with each year of age. When mortality is Gompertz, equation (6) shows that the mode satisfies

$$\mu(a_0) = \beta,$$

so the modal age at death is

$$a_0 = \frac{1}{\beta}\log(\beta/\mu_0). \tag{11}$$

These results and those following are summarized in Table 1. We expect a_0 to decrease if β increases, a property that holds for (11) so long as $a_0 > 1/\beta$ which is true for any plausible human mortality pattern. The density of age at death for the Gompertz model is

$$\phi(a) = \mu(a)e^{-(\mu(a)-\mu_0)/\beta}.$$

Because the Gompertz mortality rises exponentially, the density ϕ falls steeply at very high ages.

Equation (8) reveals a key result of this paper: the variance in adult age at death for the Gompertz model is approximately given by

$$\sigma^2 \approx \frac{1}{\beta^2}. \tag{12}$$

Thus the Gompertz variance in age at death depends only on the slope parameter β and not on μ_0. It is possible to obtain an exact expression for the variance by analytical integration in terms of special functions, but the results are not especially illuminating. However, we have computed numerically the exact variance for a range of values of β and μ_0 that are appropriate for twentieth century human mortality. We find that the exact value of σ depends only weakly on μ_0 and that equation (12) is a very accurate approximation.

It follows that a Gompertz model can only describe differences in the variance of the adult age at death with differences in the Gompertz parameter β, the age-slope of log mortality. There is a one-to-one inverse relationship between the age-slope of mortality and the variance in length of life. Were equation (12) an exact relationship, $\beta = 0.087$ would be consistent with $\sigma = 11.5$.

The Logistic Model

When measured by S_{10}, U.S. levels of σ are higher than implied by $\beta = 0.087$, more in the range of 15.0 rather than 11.5. We know there are departures from linearity in log mortality rates at advanced ages; how does this affect σ as a function of β? We write the logistic model for mortality as

$$\mu(a) = \frac{e^{\beta\, a}}{C + e^{\beta\, a}},$$

where C is the asymptote, commonly set to equal unity. Integration shows that the probability density of deaths is

$$\phi(a) = (C+1)^{1/\beta}\ \frac{e^{\beta\, a}}{(C + e^{\beta\, a})^{(1+1/\beta)}}.$$

This density falls as a simple exponential e^{-a} for high ages, much more slowly than for the Gompertz model. For the logistic, the modal age at death is

$$a_0 = \frac{1}{\beta}\, \log\,(\beta C), \tag{13}$$

and the approximate variance from equation (8) is found to be

$$\sigma^2 = \frac{(1+\beta)}{\beta^2}. \tag{14}$$

Thus the logistic also displays the remarkable property that the variance in age at death depends only on the slope parameter β. It follows that a logistic model can only describe changes in the variance of the adult age at death if the slope parameter β changes with time.

Note that if we fit a Gompertz model and a logistic model to a particular data set, the value of β must be similar in both. To see why, compare the two models near $a = 0$ which here indicates the start of adult age. With the same β, the logistic model implies only a slightly larger variance in age at adult death than the Gompertz. We expect this difference

because the density ϕ for the logistic model shallows as age increases; see the discussion after equation (9). For $\beta = 0.087$, equation (14) implies $\sigma = 12.0$, closer to reality than the Gompertz but not by much.

Of course, when we measure σ with S_{10}, we are including variance due to traffic accidents, violence, and other causes that asymmetrically impact the young in a decidedly non-Gompertz way. For the U.S., this problem may be particularly acute. Edwards and Tuljapurkar (2005) find that removing external-cause mortality reduces S_{10} by 1–1.5 years, leaving still perhaps 1.5 years in extra S_{10} that cannot be well explained by logistic or decelerating log mortality, or by external causes. A natural next step is to examine the connection between frailty in mortality and variance in length of life, two concepts that are related.

General Mortality, Multiplicative Frailty

Following Vaupel, Manton and Stallard (1979), suppose that every individual has a random frailty Z and that $g(z)\,dz$ is the probability that Z takes values between z and $z + dz$. Then mortality is determined by frailty Z and a baseline mortality function $\mu(a)$ as the product, $Z\,\mu(a)$. Conditional on frailty, the probability distribution of age at death is

$$\phi(a|Z) = Z\,\mu(a)\,\exp(-Z\,M(a)),$$

with cumulative mortality $M(a)$ defined as in equation (1). The usual specification of a frailty distribution assumes that average frailty is 1, and that the distribution of frailty has some variance $s^2 > 0$. The population probability distribution of age at death is the expectation over frailty,

$$\phi(a) = \mathcal{E}\left[\phi(a|Z)\right] = \int g(z)\,\phi(a|z)\,dz. \tag{15}$$

It is convenient to define the following averages with respect to frailty:

$$h_j(a) = \mathcal{E}\left[Z^j \; e^{-Z\,M(a)}\right], \text{ for } j = 1, 2, 3. \tag{16}$$

In the population, the modal age at death must then satisfy

$$\mu'(a) = \mu^2(a) \; \frac{h_2}{h_1}. \tag{17}$$

Note that if every frailty were equal to 1, we would have $h_2 = h_1$ and this equation would reduce to our earlier equation (6). To approximate the variance in age at death, we use equation (8) and obtain

$$\phi''(a_0) = h_1\,\mu'' + \mu^3 \left\{ h_3 - 3\left(\frac{h_2^2}{h_1}\right)\right\}, \tag{18}$$

where the h_i are evaluated at the mode a_0, and then set

$$\sigma^2 = \frac{\phi(a_0)}{|\phi''(a_0)|} = \frac{h_1(a_0)\mu(a_0)}{|\phi''(a_0)|}.$$

In these expressions, if every frailty is set equal to 1, we have $h_3 = h_1$ and the variance σ^2 reduces to the value in equation (8).

General Mortality, Gamma Multiplicative Frailty

The expressions we presented thus far are not as illuminating as one might hope about how frailty would affect the mode or the variance in age at death . To obtain a qualitative sense of the effect of frailty, we consider the case when frailty Z follows a gamma distribution (Vaupel, Manton and Stallard, 1979). In this case, the probability that Z lies between w

and $w + dw$ is assumed to be $g(w)\,dw$ with

$$g(w) = \frac{k^k}{\Gamma(k)}\, w^{k-1}\, e^{-k}, \tag{19}$$

where the average frailty is 1 and the variance of frailty is $\mathrm{Var}(Z) \equiv s^2 = 1/k$. This distribution is convenient, as Vaupel et al. pointed out, because we can use it with any baseline mortality $\mu(a)$ to find an explicit expression for the population average distribution of age at death, whose general form was given by equation (15):

$$\phi(a) = \mu(a) \left(\frac{k}{(k + M(a))} \right)^{k+1}. \tag{20}$$

We can differentiate to find that the modal age at death is defined by the condition

$$\mu'(a_0) = \left(\frac{1 + s^2}{1 + s^2 M(a_0)} \right) \mu^2(a_0). \tag{21}$$

Notice that if all individuals had the same frailty, $s^2 = 0$ and equation (21) would reduce to the simpler equation (6). Qualitatively, the denominator on the right describes how frailty alters the rate of change of average mortality and survival depending on how much selection acts against more frail individuals. The magnitude of selection depends on both the variance s^2 in frailty, and the cumulative mortality hazard $M(a)$. Strong selection will act to decrease the modal age at death.

We can find the second derivative ϕ'' at the modal age and use it in a Taylor approximation similar to equation (7) in order to obtain the variance in age at death:

$$\begin{aligned} \sigma^2 &= \left(\frac{\phi(a_0)}{|\phi''(a_0)|} \right), \\ &= \frac{\mu(a_0)}{|\mu''(a_0) - \mu^3(a_0)\ [(1 + s^2)(2 + s^2)] \,/\, [(1 + s^2 M(a_0)^2)]\,|}. \end{aligned} \tag{22}$$

Note again the selection effect in the denominator in which s^2 is multiplied by the cumulative

mortality $M(a_0)$ which has occurred at ages below the mode. Strong selection via a large $M(a_0)$ will combine with variance in frailty s^2 to reduce the denominator and thus inflate the variance σ^2 in age at death.

Gompertz Mortality, Gamma Multiplicative Frailty

We can learn more by combining a Gompertz baseline mortality $\mu(a) = \mu_0\, e^{\beta\, a}$ with multiplicative gamma-distributed frailty. The modal age at death for this model is found using equation (21) with the Gompertz mortality, and yields the condition

$$\mu(a_0) = (\beta - s^2\mu_0),$$

which explicitly gives us the mode as

$$a_0 = \frac{1}{\beta} \log\left(\beta/\mu_0 - s^2\right). \tag{23}$$

Compared with equation (11) for the standard Gompertz, which is shown in the first column, third row of Table 1, this equation reveals how frailty acts to reduce the modal age at death.

The variance in age at death is obtained using equations (21) and (22) and a little algebra, which yields the remarkably simple result that

$$\sigma^2 = \frac{(1 + s^2)}{\beta^2}. \tag{24}$$

Comparing this with equation (12) for the standard Gompertz, both shown along the bottom row in Table 1, shows that frailty amplifies the variance in age at death. A Gompertz model with gamma frailty can describe changes in the variance of the adult age at death with changes either in the Gompertz slope parameter β or in the variance s^2 in frailty or in both.

Summary

We have arrived at a set of important positive results. A key finding is that the Gompertz slope of log mortality through age is inversely related to the variance in length of life. Differences across time, space, or SES in the latter can only be captured by differences in age-slopes of mortality. This insight is not greatly altered if mortality follows a logistic curve, flattening out at advanced ages. Injecting frailty into standard mortality models loosens the relationship between variance in life span and the Gompertz slope, often called the rate of aging. A Gompertz model with Gamma multiplicative frailty allows us to model heterogeneity in variance as deriving from heterogeneity in either the Gompertz slope, in the variance in frailty, or in both.

But we view any normative or etiological insights from our exploration of these models as considerably more elusive. On the one hand, we think the basic result concerning variance in the Gompertz slope is important. Some have viewed the Gompertz slope as a kind of species-specific parameter that is eternally fixed while the intercept may fluctuate (Finch, Pike and Witten, 1990). But we show it can only be constant if we add a free parameter like frailty and allow it to fluctuate arbitrarily in order to fit the data. A more straightforward reading of the evidence is that instead, both nature and nurture must affect the Gompertz slope, at least in human populations and potentially in other species.

The addition of multiplicative frailty to a Gompertz model only addresses changing variances in age at death if we assume that frailty distributions have been changing quite rapidly over time. Temporal change in frailty has not been a feature of mortality models, and it is not clear why the distribution of such frailties would narrow over time. From an evolutionary perspective, it is not altogether clear why frailty should have persisted into modern humans at all. It is also unclear why in the cross section African Americans should have persistently higher frailty than whites, or why Americans in general should endure higher levels than Europeans or Japanese. Frailty models may allow us to model mortality better in a strictly mechanical sense, but they do not appreciably improve our understanding.

Models of Adult Mortality in Practice

These results concerning variance in length of life and the age-slope of mortality bear implications for modeling and forecasting. In this section we examine how variance is implicitly or explicitly treated in an array of common frameworks. We first examine forecasting models, including a recent technique that involves mortality "translation," which we explain below, as well as the the popular Lee and Carter (1992) forecasting model, which can be seen as a generalized Gompertz model. Then we discuss several standard mortality models that are commonly used in short panels with microdata, where cross-sectional patterns in variance are more important. We examine the Cox (1972) model, the class of nonparametric models suggested by the methods of Kaplan and Meier (1958), and accelerated failure time (AFT) models, which some have associated with frailty models.

Aggregate Forecasting Models

Mortality Translation: Bongaarts-Feeney

The pioneering model of mortality translation due to Bongaarts and Feeney (2003) provides an appealingly simple description of mortality change, although it is controversial. A recent edited volume by Barbi, Bongaarts and Vaupel (2008) provides a thorough overview of the method, which is related to the concept of tempo effects in mortality (Bongaarts and Feeney, 2002).

One can describe the Bongaarts-Feeney model in terms of a hypothetical cohort following a period life table. Let T_1 be the random age at death of an individual in this cohort in period t_1. In a later period $t_2 > t_1$, suppose that the effect of mortality change between the two periods is completely described by delaying each death by the same amount. We assume infant mortality is practically zero and thus can also be delayed in this fashion, which although unrealistic is consistent with our focus on adult mortality. Each random age at death T_1 in the first period is thus replaced in the later period by the random age at death

$T_1 + D$, where $D > 0$ is fixed. We see at once that the average age at death increases from $e_0(t_1) = \mathcal{E}[T_1]$ in period t_1 to $e_0(t_2) = e_0(t_1) + D$ in the later period t_2. If we increase the mean age at death by some fixed annual amount, we have found a model of mortality change that describes a constant trend in e_0. We use the term mortality translation for any such model.

Notice that translation only affects the mean age at death and not its variance. Shifting every random age at death from T_1 to $T_2 = T_1 + D$ results in a constant variance, $\mathrm{Var}(T_1) = \mathrm{Var}(T_2)$, so long as D is fixed. In fact, translation leaves unchanged all the central moments of the random age at death. Put geometrically, translation necessarily implies that the shape of the distribution of age death does not change.

Mortality translation is appealing because it can be used with any mortality model. Bongaarts and Feeney (2003) and Bongaarts (2005) used translation for a Gompertz and a logistic model. Vaupel (1986) used a Gompertz model in essentially the same way, although he did not explicitly refer to translation. Take any reasonable adult mortality function $\mu(a)$. In period t_1 suppose that the corresponding mortality $\mu_1 > 0$ for ages greater than some cutoff age A_1. In a later period t_2 define adult mortality to be a translation of the original mortality schedule:

$$\mu_2(a) = \begin{cases} 0 & \text{if} \quad A_1 \leq a < (A_1 + D) \\ \mu_1(a - D) & \text{if} \quad a \geq (A_1 + D) \end{cases}$$

It follows automatically that the probability distribution of ages at death is also translated. If ϕ_1 and ϕ_2 are the distributions in the two periods, then

$$\phi_2(a) = \phi_1(a - D), \text{ for } a \geq (A_1 + D).$$

This is simply an alternative statement of the translation of the random age at death T_1 distributed as ϕ_1 to $T_1 + D$.

It is obvious that mortality translation, by construction, cannot describe temporal changes in the variance in the probability distribution of age at death, or for that matter, of other central moments of ϕ related to the shape of the distribution, such as skewness or kurtosis. As revealed by Wilmoth and Horiuchi (1999), Cheung et al. (2005), and Edwards and Tuljapurkar (2005), trends in all of these moments have been and remain very interesting even in industrialized countries. Prior to 1960, S_{10} was strongly declining in industrialized countries, a pattern that is repeating itself among many developing countries today (Edwards, 2009*c*). To be sure, Bongaarts and Feeney (2003) never intended their model to apply universally across historical periods; their aim was mortality forecasting for the U.S. and other industrialized countries. But while S_{10} has remained roughly steady on average in the U.S. since 1960, it has also fluctuated up and down within a 1.5-year band over time (Edwards and Tuljapurkar, 2005), or $\pm 5\%$. Forecasting via translation implicitly makes a relatively bold statement about long-term trends in the variance as well as the tolerance for error in estimating it.

Generalized Gompertz: Lee-Carter

Lee and Carter (Lee and Carter, 1992) proposed a parsimonious three-parameter model that explains temporal trends in mortality well in industrialized countries (Lee and Miller, 2001). Using the singular value decomposition, they estimate

$$\log \mu(a,t) = \mathbf{a}(a) + \mathbf{b}(a)\ k(t), \tag{25}$$

where $\mu(a,t)$ is the mortality at age a in period t, $\mathbf{a}$ and $\mathbf{b}$ are constant age profiles or vectors, and $k(t)$ is a random walk with negative drift.

The intercept vector $\mathbf{a}$ is an average of age-specific log mortality rates over the historical sample period, so it ends up being approximately Gompertz or logistic in shape. But the $\mathbf{b}$ vector is not necessarily constant with age, as it would have to be in a Gompertz model

with a fixed age slope over time. Indeed, fits of **b** typically reveal stark differences in rates of mortality decline across age in industrialized countries (Tuljapurkar, Li and Boe, 2000). By consequence, the slope and curvature of mortality in this model are free to evolve over time, which as our results show can easily lead to changes in the variance of age at death. The singular value decomposition of equation (25) produces optimal fits of age-specific mortality rates; the moments of the distribution of ages at death are backed out of the model rather than hard-wired as in the Bongaarts-Feeney translation model.

Figure 4 depicts historical data since 1959 and probabilistic forecasts of S_{10} for U.S. out to 2050 using the Lee-Carter model applied to log mortality rates for both sexes combined from the Human Mortality Database (2009). The forecast is shown by three dashed lines that indicate the 2.5, 50, and 97.5 percentiles of the distribution. Several characteristics are informative. There is clearly an extrapolated trend in S_{10}, which declines in the median forecast from 15 years around 2000 to just above 14 years by 2050. While this looks like a bold assumption in the figure, by comparison practically all high-income countries except for France already enjoy levels of S_{10} around 14 or below. A second point is that although the forecast confidence interval is nearly 1 year wide by 2050, it does not appear to allow for the same kind of volatile fluctuations we see in the time series for S_{10}. Like e_0, S_{10} is a smoothed average of the many age-specific mortality rates that the Lee-Carter model actually forecasts. It is not surprising that forecasts of S_{10}, like e_0, would be considerably less volatile than historical data.

Implications

Over the last two centuries, the variance in age at adult death as measured by the standard deviation S_{10} has declined by almost 50%. Were we to use a Gompertz model to describe period mortality at ages over 10, then the slope of the Gompertz model would have to increase by about 40% in order to replicate observed trends in S_{10}. A logistic model for period mortality would require a larger increase, about 50%, in the slope. Over the past

50 years, the trend in S_{10} has been toward much more gradual decline, perhaps 1% in the U.S., but temporary fluctuations in the variance around its slight downward trend have been much larger, more like 10%.

Because mortality translation models do not allow for any change in the variance of adult death, we doubt they should be used to say anything about the age pattern of deaths, even if they may describe changes in e_0 perfectly well. A striking result about long run mortality change is the demonstration by Oeppen and Vaupel (2002) that world-record high e_0 has risen at a remarkable linear rate over the past 160 years. Such uniformity suggests that upper bounds on life span, prognosticated and then consistently broken throughout this period, are not as clear as some currently believe, and that the pace of human development and achievement measured in this way has been rapid and surprisingly steady across several distinct periods of socioeconomic and epidemiological transitions. But Edwards and Tuljapurkar (2005) describe how long-term trends in the record-low variance paint a very different picture regarding the gains in human well-being along the dimension of mortality. True, progress against the Gompertz slope has indeed been achieved, contrary to the opinions of those who may have viewed it as immutable. But long-term gains have come more in fits and starts rather than continuously, and this highlights the remaining challenges, as does the considerable heterogeneity across countries this century in progress against variance. We do not fully understand the sources of variance in life spans, nor the underlying health inequalities they presumably reflect, and this is a problem for policy as well as for modeling and forecasting mortality. Still, we acknowledge that translation models nonetheless provide a good statistical fit to mortality patterns in industrialized countries since 1950 (Bongaarts, 2005). But the shape of the age distribution of deaths provides enough important insights that we should probably not assume it to be fixed over long periods of time.

By contrast, the Lee and Carter (1992) forecasting model can and does predict changes in the variance in the age at death, but with caveats. Its implicit forecasts of S_{10} appear to be smooth extrapolations of average trends over the forecast interval. This is consistent with

the spirit and mechanics of the Lee-Carter framework and mirrors patterns in Lee-Carter forecasts of e_0. But since the historical paths of S_{10} have been considerably more convoluted than those of e_0 in every industrialized country, the simple extrapolation of the long-term average trend in S_{10} produced by Lee-Carter seems incongruous. We conclude that while the flexibility of Lee-Carter makes it a valuable model of mortality, it does not do as good a job predicting trends in variance as we believe a model probably should.

Short Panel Models

Cox Proportional Hazards

The semiparametric model of Cox (1972) is often said to be the gold standard for modeling survival in clinical settings and other relatively short panels. The model's lone assumption is that hazards are proportional across groups identified by covariates. Meanwhile, the shape of the underlying mortality mortality is not parameterized.

Although the Cox model does not require background mortality to be Gompertz, the fact remains it often is, at least approximately, because mortality rates tend to increase exponentially until very advanced ages. Given this, our theoretical result that ties the Gompertz slope to the variance in length of life also implies that the Cox proportional hazards assumption is flatly inconsistent with cross-sectional patterns in variance. In addition to having lower mean length of life, low-status groups also have higher variance, which means they must have a smaller Gompertz slope or less steeply increasing mortality with age. Thus it is clear that hazards are not proportional across individuals of varying ages and statuses. By incorrectly assuming that they are, the standard Cox model will fail to model any differences in subgroup variance.

It is by no means a revelation that hazards are often not proportional across groups, of course. Researchers have developed a variety of tools to test for violations of proportionality and a set of alternative models when violations are found (Hess, 1995; Lee and Go, 1997; Therneau and Grambsch, 2000). One solution is to use time-varying coefficients; another is to

perform stratified Cox regression. In either case, the analyst must judge which characteristics on which to stratify or to allow to vary over time. Our results suggest that allowing the effect of race or SES to vary systematically over time should adequately capture the different subgroup variances and Gompertz slopes. Whether it is sufficient to use just one of these covariates rather than several to model differences in the age-slope is unknown.

Accelerated Failure Time Models

Another class of tools that researchers use when the proportional hazards assumption is violated are accelerated failure time (AFT) models (Therneau and Grambsch, 2000). These are parametric models of survival time rather than of the hazard, which take the form

$$\log T = a + XB + \epsilon \tag{26}$$

where X is a vector of covariates and ϵ is an error term. Interestingly, AFT models have been discussed in the context of modeling mortality with multiplicative frailty (Hougaard, 1995), as well as in an alternative conceptualization of tempo effects in mortality (Rodríguez, 2008).

But a proportional scaling of length of life between racial subgroups or across SES, such as we find in the AFT model, is even more misspecified than a proportional scaling of hazards vis-à-vis the variance. While any scaling, additive or proportional, can provide adequate fit to the mean lengths of life across subgroups, a proportional scaling implies that the error in the level of length of life, namely the variance, becomes larger with longer mean length. This is precisely the opposite of what we see in data.

While it is tempting to suggest a linear modeling of survival time as an alternative, the error term in such a model would clearly be heteroscedastic. Without correction, the robustness of inference would suffer, and like the standard Cox, a linear AFT model would fail to model subgroup variances correctly.

Kaplan-Meier Estimation

Another alternative to the Cox model is the fully nonparametric estimator of the survivorship function due to Kaplan and Meier (1958). Given the difficulties standard parametric models seem to face in capturing the variance in length of life, which is reflected in the slope of survivorship, a nonparametric approach would seem at first to be ideal. One challenge is that the Kaplan-Meier estimator may sacrifice efficiency when compared with parametric approaches (Miller, Jr., 1983). Another challenge is that identifying which covariates ought to matter and how is left entirely up to the researcher.

Implications

Subgroup differences in variance in length of life and the varying age-slopes in mortality rates they imply are a large problem for short panel models. Without correction, the Cox proportional hazards model will probably produce biased results if age is correlated with any of the other covariates, and it will definitely fail to capture subgroup differences in life-span variance. Modeling time-varying covariates, which are really age-varying and thus exactly what is needed, is one solution, and stratified Cox regression is another, but both will reduce power. While AFT models appear if anything to worsen the modeling problem, the Kaplan-Meier estimator would clearly improve it. The question is whether the efficiency cost exceeds the costs of misspecification or reduced power, and the exact tradeoff is likely to be application-specific. We intend this section merely as a renewed warning, with new theoretical underpinnings, about a known but important issue in micro-level modeling.

Summary and Discussion

Our primary message is that the Gompertz slope, the rate of increase in mortality through age, cannot be constant across time, space, or characteristic, because neither is the variance in length of life. The Gompertz slope, which is often conceptualized as the rate of aging,

typically increases with higher status, although aggregate patterns across OECD countries do not neatly fit this simplified view. The result of a steeper Gompertz slope is a reduction in the variance around length of life, and it always seems to be accompanied by a reduction in the Gompertz intercept, which raises life expectancy.

Variance in length of life is costly, whether viewed at the population level as aggregate health inequality or at the individual level as a mean-preserving spread in how long we live (Edwards, 2009*b*). When rephrased as a faster rate of aging, a steeper Gompertz slope may not sound like something good. But the compression of mortality around an ever-increasing adult mode age at death represents decreases in costly uncertainty and arguably preventable premature death.

Historical data reveal a massive decrease in the uncertainty around adult length of life concomitant with revolutions in nutrition, public health, and the prevention of communicable disease. Although progress against variance in developing countries has continued (Edwards, 2009*c*), compression has largely stalled in industrialized countries since 1960 (Wilmoth and Horiuchi, 1999; Edwards and Tuljapurkar, 2005). As we discussed, a critical question for forecasting models is how future patterns in the variance of age at death will unfold.

Given the long-term nonstationarity of variance, and the rapidity of its decline prior to 1960, it is questionable how helpful extrapolative forecasts may be. This seems especially true given the pace of scientific advancement in the genetics of aging, an entirely new field with the potential to change the way we age. In a recent review, de Magalhães, Cabral and Magalhães (2005) discuss how a subset of genetic interventions in lab mice appear to have steepened the Gompertz slope, compressing mortality and "decelerating" aging. At the same time, other interventions do not. It is an unanswered but provocative question whether such new knowledge may be able to rectangularize human survivorship even further in a future gerontological revolution.

This is not to say that gains against the diseases of old age are the only possible source of gains against variance, which is almost certainly not true. That variance is higher among

subgroups with lower SES is a fairly clear indication that we can make considerable progress against variance by addressing the disease of socioeconomic inequality. While there is no clear link between population S_{10} and income or educational inequality, it is telling that we see racial differences in S_{10}, which could in principle be driven by genetic differences, that look almost identical to SES differences. This is a straightforward extension of the conclusion that racial differences in mortality seem to be a manifestation of racial differences in SES (Preston and Taubman, 1994). But reducing inequality in SES is likely to be a function of political will, which few demographers and health economists would be comfortable forecasting.

References

Arias, Elizabeth. 2007. "United States Life Tables, 2004." *National Vital Statistics Reports* 56(9).

Barbi, Elisabetta, John Bongaarts and James W. Vaupel, eds. 2008. *How Long Do We Live?* New York: Springer.

Bongaarts, John. 2005. "Long-Range Trends in Adult Mortality: Models and Projection Methods." *Demography* 42(1):23–49.

Bongaarts, John and Griffith Feeney. 2002. "How Long Do We Live?" *Population and Development Review* 28(1):13–29.

Bongaarts, John and Griffith Feeney. 2003. "Estimating mean lifetime." *Proceedings of the National Academy of Sciences USA* 100(23):13127–13133.

Caldwell, John C. 1976. "Toward A Restatement of Demographic Transition Theory." *Population and Development Review* 2(3/4):321–366.

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.

Cox, D. R. 1972. "Regression Models and Life-Tables." *Journal of the Royal Statistical Society. Series B (Methodological)* 34(2):187–220.

de Magalhães, João Pedro, José A. S. Cabral and Domingo Magalhães. 2005. "The Influence of Genes on the Aging Process of Mice: A Statistical Assessment of the Genetics of Aging." *Genetics* 169(1):265–274.

Edwards, Ryan D. 2009*a*. "The Cost of Cyclical Mortality." *B.E. Journal of Macroeconomics,* 9(1):Contributions, Article 7.

Edwards, Ryan D. 2009*b*. "The Cost of Uncertain Life Span." Unpublished manuscript, March.

Edwards, Ryan D. 2009*c*. "Trends in World Inequality in Life Span Since 1970." Unpublished manuscript, May.

Edwards, Ryan D. and Shripad Tuljapurkar. 2005. "Inequality in Life Spans and a New Perspective on Mortality Convergence Across Industrialized Countries." *Population and Development Review* 31(4):645–675.

Finch, Caleb E., Malcolm C. Pike and Matthew Witten. 1990. "Slow Mortality Rate Accelerations During Aging in Some Animals Approximate That of Humans." *Science* 249(4971):902–905.

Fries, James F. 1990. "Aging, Natural Death, and the Compression of Morbidity." *New England Journal of Medicine* 303(3):130–135.

Glei, Dana A. and Shiro Horiuchi. 2007. "The narrowing sex differential in life expectancy in high-income populations: Effects of differences in the age pattern of mortality." *Population Studies* 61(2):141–159.

Gompertz, Benjamin. 1825. "On the nature of the function expressive of the law of human mortality and on a new mode of determining life contingencies." *Philosophical Transactions of the Royal Society of London* 115:513–585.

Hess, Kenneth R. 1995. "Graphical Methods for Assessing Violations of the Proportional Hazards Assumption in Cox Regression." *Statistics in Medicine* 14(15):1707–1723.

Hougaard, Philip. 1995. "Frailty Models for Survival Data." *Lifetime Data Analysis* 1(3):255–273.

Human Mortality Database. 2009. University of California, Berkeley (USA) and Max Planck Institute for Demographic Research (Germany). Available at www.mortality.org.

Kannisto, Väinö. 2000. "Measuring the Compression of Mortality." *Demographic Research* 3(6).

Kaplan, Edward L. and Paul Meier. 1958. "Nonparametric Estimation from Incomplete Observations." *Journal of the American Statistical Association* 53(282):457–481.

Lee, Elisa T. and Oscar T. Go. 1997. "Survival Analysis in Public Health Research." *Annual Review of Public Health* 18:105–134.

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

Lee, Ronald D. and Timothy Miller. 2001. "Evaluating the Performance of the Lee-Carter Approach to Modeling and Forecasting Mortality." *Demography* 38(4):537–549.

Miller, Jr., Rupert G. 1983. "What Price Kaplan-Meier?" *Biometrics* 39(4):1077–1081.

Oeppen, Jim and James W. Vaupel. 2002. "Broken Limits to Life Expectancy." *Science* 296:1029–1031.

Olshansky, S. Jay, Bruce A. Carnes and Christine Cassel. 1990. "In Search of Methuselah: Estimating the Upper Limits to Human Longevity." *Science* 250(4981):634–640.

Preston, Samuel H. 1975. "The Changing Relationship between Mortality and Level of Economic Development." *Population Studies* 29(2):231–248.

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.

Rodríguez, Germán. 2008. Demographic Translation and Tempo Effects: An Accelerated Failure Time Perspective. In *How Long Do We Live?*, ed. Elisabetta Barbi, John Bongaarts and James W. Vaupel. New York: Springer pp. 69–92.

Thatcher, A. Roger, Väinö Kannisto and James W. Vaupel. 1998. The Force of Mortality at Ages 80 to 120. In *Monographs on Population Aging*. Number 5 Odense, Denmark: Odense University Press.

Therneau, Terry M. and Patricia M. Grambsch. 2000. *Modeling Survival Data: Extending the Cox Model*. New York: Springer.

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.

Tuljapurkar, Shripad, Nan Li and Carl Boe. 2000. "A universal pattern of mortality decline in the G7 countries." *Nature* 45:789–792.

Vaupel, James W. 1986. "How Change in Age-specific Mortality Affects Life Expectancy." *Population Studies* 40(1):147–157.

Vaupel, James W., James R. Carey, Kaare Christensen, Thomas E. Johnson, Anatoli I. Yashin, Niels V. Holm, Ivan A. Iachine, Väinö Kannisto, Aziz A. Khazaeli, Pablo Liedo,

Valter D. Longo, Yi Zheng, Kenneth G. Manton and James W. Curtsinger. 1998. "Biodemographic Trajectories of Longevity." *Science* 280(5365):855–860.

Vaupel, James W., Kenneth G. Manton and Eric Stallard. 1979. "The Impact of Heterogeneity in Individual Frailty on the Dynamics of Mortality." *Demography* 16(3):439–454.

Wachter, Kenneth W. and Caleb E. Finch, eds. 1997. *Between Zeus and the Salmon: The Biodemography of Longevity*. Washington, DC: National Academy Press.

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

Wilmoth, John R. and Shiro Horiuchi. 1999. "Rectangularization Revisited: Variability of Age at Death within Human Populations." *Demography* 36(4):475–495.

Table 1: Characteristics of the distribution of adult death in three models

Parameter	**Gompertz**	**Logistic**	**Gompertz Gamma**
Mortality at age a, $\mu(a)$	$\mu_0 e^{\beta a}$	$\frac{e^{\beta a}}{C + e^{\beta a}}$	$Z\mu_0 e^{\beta a}$
Density of deaths at a, $\phi(a)$	$\mu(a) e^{-(\mu(a)-\mu_0)/\beta}$	$(C+1)^{1/\beta} \frac{e^{\beta a}}{(C + e^{\beta a})^{1+1/\beta}}$	$\mu(a) \left(\frac{k}{k + M(a)} \right)^{k+1}$
Mode age at death, a_0	$\frac{1}{\beta} \log(\beta/\mu_0)$	$\frac{1}{\beta} \log(\beta C)$	$\frac{1}{\beta} \log(\beta/\mu_0 - s^2)$
Variance in age at death, σ^2	$\frac{1}{\beta^2}$	$\frac{1+\beta}{\beta^2}$	$\frac{1+s^2}{\beta^2}$

Notes: In the logistic model, C is the asymptote, commonly set to equal one. The cumulative force of mortality at age a is $M(a) = \int_0^a ds\ \mu(s)$. In the Gompertz Gamma model, the multiplicative frailty index Z is distributed gamma with density equal to $g(w) = \frac{k^k}{\Gamma(k)} w^{k-1} e^{-k}$, an average $\mathcal{E}(Z) = 1$, and a variance $\text{Var}(Z) \equiv s^2 = 1/k$.

Figure 1: Distributions of age at death by race in the U.S. in 2004

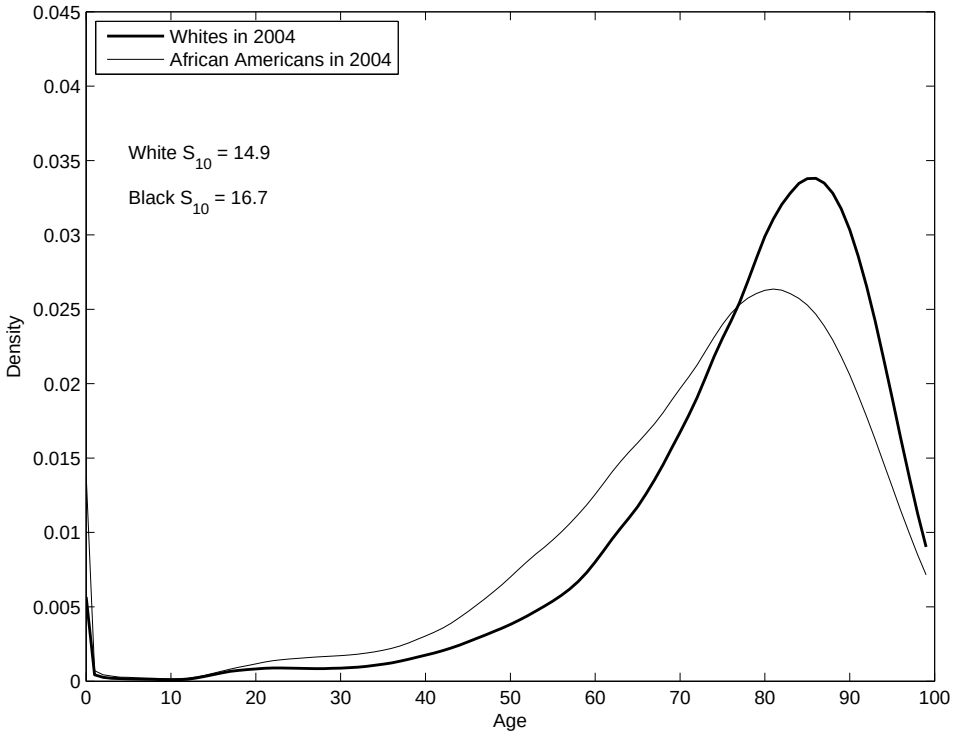

Source: Arias (2007). Data are life-table deaths by race ${}_nd_x$ for both sexes combined.

Figure 2: Distributions of age at death in Sweden in 1900 and 2000

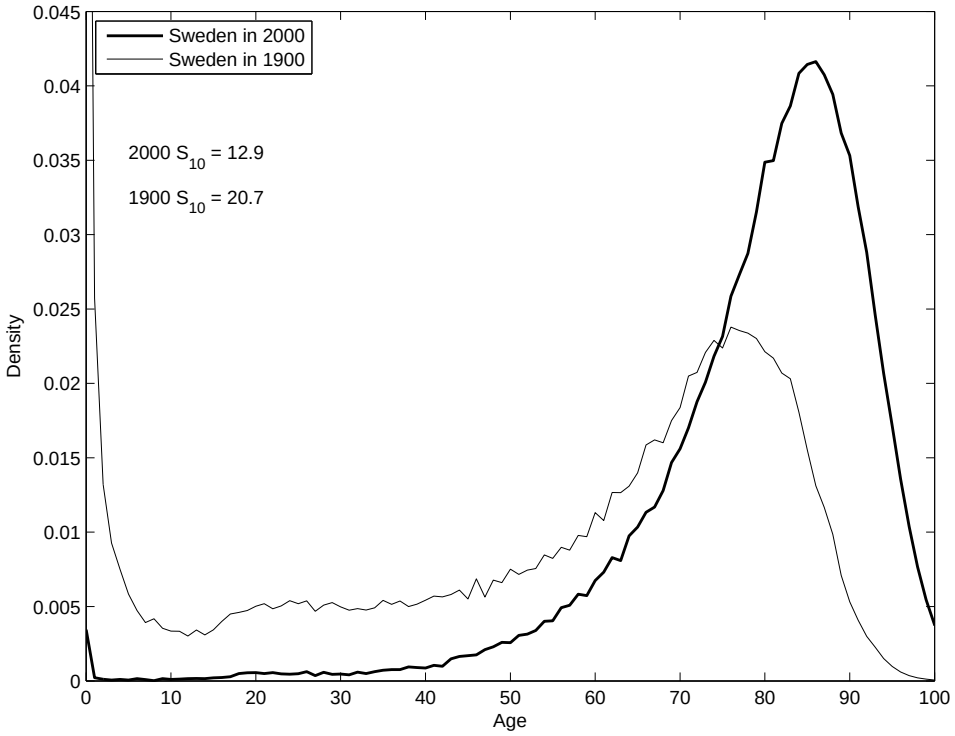

Source: Human Mortality Database (2009). Data are life-table deaths ${}_nd_x$ for both sexes combined.

Figure 3: Decomposing total variance in length of life in Sweden since 1900

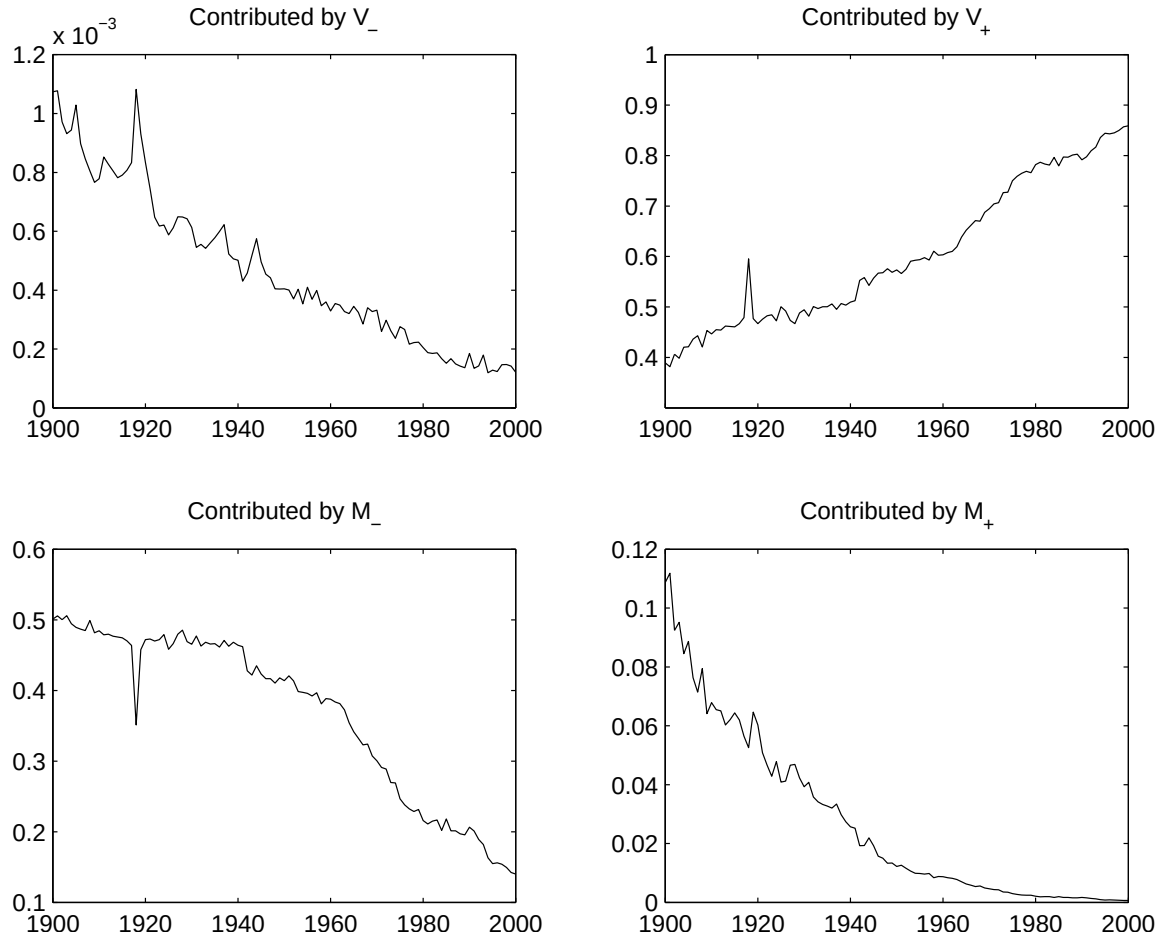

Source: Human Mortality Database (2009) and authors' calculations. Each panel shows a component's proportion of total variance in period length of life, which is the weighted sum of squared differences between life expectancy at birth, e_0, and midpoints of single years of age up to 110, where the weights are life-table deaths ${}_nd_x$. See equation (5) in the text for the formulation of each component; V_- is the variance of age at death below age 10 weighted by deaths below age 10, V_+ is the variance above age 10 weighted by deaths above 10, M_- is the squared difference between e_0 and the mean age at death below age 10 weighted by deaths below 10, and M_+ is the squared difference between e_0 and the mean age at death above 10 weighted by deaths above 10.

Figure 4: Historical patterns and Lee-Carter forecasts of S_{10} for both sexes combined in the U.S.

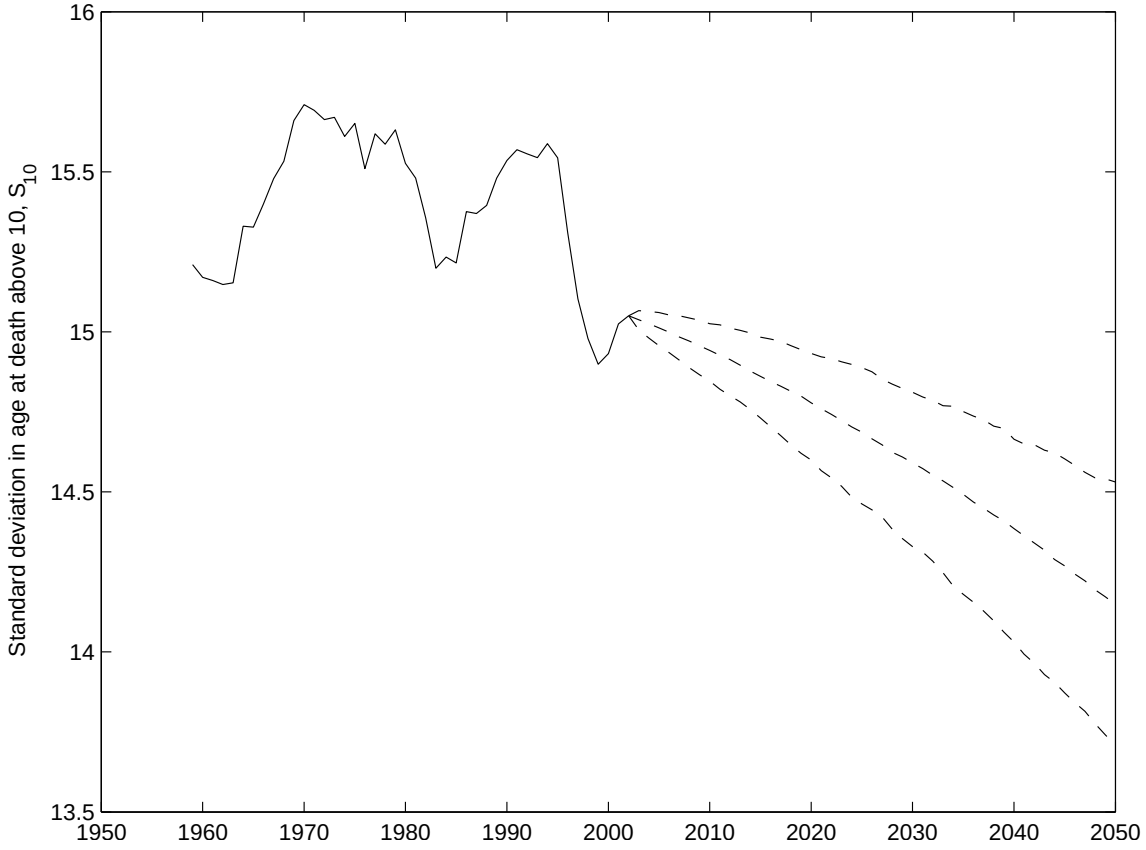

Source: Human Mortality Database (2009) and authors' calculations. S_{10} is the standard deviation in ages at death based on the period life table. The authors fit the Lee-Carter model as shown in equation (25) via singular value decomposition applied to log age-specific mortality rates for both sexes combined below age 100 from the HMD. Dashed lines represent the 2.5, 50, and 97.5 percentiles of the probabilistic forecast made assuming $k(t)$ follows a random walk with drift.