Mortality Deceleration and Mortality Selection:

Three unexpected implications of a simple model

Elizabeth Wrigley-Field

wrigleyfield@wisc.edu

ABSTRACT:

Mortality deceleration is commonly attributed to mortality selection: as a cohort ages, its frailest members die, potentially leading aggregate mortality to decelerate. Drawing on formal analysis and mortality simulations, I dispute some common assumptions about how population heterogeneity produces deceleration. I show that even in a very simple model—one composed of just two subpopulations with Gompertz mortality— (1) aggregate mortality can decelerate even while a majority of the population is frail; (2) multiple decelerations are possible; and (3) mortality selection can produce acceleration as well as deceleration. Simulations show that these patterns are plausible in model populations that in the aggregate resemble those in the Human Mortality Database. I argue that these results: challenge some conventional heuristics for understanding the relationship between selection and deceleration; suggest that standard parametric models, assumed to plateau at most once, may sometimes mislead analysts as to deceleration timing; and make problematic certain inferences from deceleration timing to patterns of social inequality.

Mortality deceleration is of deep interest to demographers hoping to understand how population heterogeneity manifests itself in population-level mortality. In the words of Lynch et al. (2003: 462), "measuring deceleration, compression and crossover is *the* means by which to examine heterogeneity within and between populations."

Deceleration—the slowing of mortality's rise with age—is considered telling evidence of unobserved heterogeneity because it is taken to arise from mortality selection: as a cohort ages, its frailest members die first, so that the remaining group increasingly is whittled to those most robust to death (Beard 1959, 1971; Vaupel, Manton and Stallard 1979; Vaupel and Yashin 1985). Mortality deceleration has been widely documented empirically (Bebbington, Lai, and Zitikis 2007; Gavrilov and Gavrilova 1991; Horiuchi and Wilmoth 1998; Lynch and Brown 2001; Lynch, Brown, and Harmsen 2003; Olshansky 1998; Vaupel 1997).

Underlying this interest in mortality deceleration is the idea that population heterogeneity and aggregate deleceration patterns are related to one another in qualitatively straightforward ways, so that one can reason from one to the other. Previous studies (Gampe 2010; Horiuchi and Wilmoth 1998; Lynch and Brown 2001; Lynch et al. 2003) have reasoned in both directions, using assumptions about unobserved heterogeneity to test mortality selection as an explanation of observed deceleration patterns, and using observed deceleration patterns to generate new hypotheses about unobserved heterogeneity. As I will show, such reasoning sometimes seems to be based on a simple heuristic: mortality decelerates when the percent frail in a cohort reaches some critically low threshold. Moreover, the analytical methods used to study mortality deceleration, and to relate deceleration to social inequality, commonly assume that cohort mortality decelerates at most once, and that mortality selection contributes only deceleration—never acceleration—to cohorts' aggregate mortality.

In this paper, I dispute these assumptions about how population heterogeneity produces deceleration. In their place, I offer three complicating analytic results that arise even in a very basic two-subpopulation model. First, aggregate mortality can decelerate even while a majority of the cohort is frail. Second, a cohort can decelerate more than once. Third, mortality selection itself, regardless of the acceleration in subpopulation mortality, can produce both deceleration and acceleration.

I argue that these results have three significant implications: they challenge some widespread intuitions about the relationship between selection and deceleration; they suggest that standard parametric models, assumed to plateau at most once, may sometimes mislead analysts as to deceleration timing; and they make problematic certain inferences from deceleration timing to patterns of social inequality.

To demonstrate these claims, I proceed in three steps. First, I present a basic two-group model and analyze mortality deceleration in terms of the increase in subpopulation mortality and the rate of selection. This formal analysis, which extends the analysis given in Vaupel and Zhang (2010), promotes an intuitive understanding of how the counter-intuitive cohort mortality patterns described here can arise. The second stage of the analysis presents simulations with parameters from the Human Mortality Database to demonstrate that these patterns can indeed arise with realistic human mortality parameters. Finally, I demonstrate the implications of these results for selection theory, for parameterizing mortality in empirical research, and for reasoning between mortality deceleration and inequality between populations. Together, the results show that even the simplest model can generate consequences so complicated as to invalidate common intuitions and inferences about mortality deceleration.

Mortality deceleration and mortality selection

Mortality deceleration

Mortality deceleration is the label given to a class of mortality patterns deviating from the exponential mortality of the Gompertz model, which posits that mortality accelerates at increasing speed as a cohort ages. In operationalizing deceleration, demographers have variously highlighted different degrees of deviation from exponential mortality. Some focus on what I call *relative deceleration*, which occurs when mortality continues to accelerate, but does so more slowly than at younger ages. This begins when the third derivative (jerk) of aggregate mortality becomes negative, or, equivalently, when the second derivative begins to decline (Rau et al. [2009] reviews the alternatives and advocates this measure). Others (Bebbington et al. 2007, Lynch and Brown 2001; Lynch et al. 2003) employ what I call *absolute deceleration*, which occurs when mortality is no longer accelerating at all: the second derivative (acceleration) is negative, and the first derivative has begun to decline.[1] In principle, cohorts may evince an even more extreme deviation from Gompertz mortality, *mortality decline*, which occurs when the first derivative (slope) is negative. Table 1 summarizes these measures. Here, I conceive of mortality deceleration as a process that begins with relative deceleration and may progress through absolute deceleration and even mortality decline, and which also may stop or reverse at any point. Accordingly, the crucial measures defining deceleration and reacceleration in what follows

[1] Lynch and Brown (2001) use the term *absolute deceleration* as I use it, but use *relative deceleration* to refer to the Lifetable Aging Rate (LAR), discussed below. They do not discuss what I call relative deceleration, which entered the demographic literature more recently, with Rau et al. (2009).

will be the signs of the first, and, particularly, the second and third derivatives of mortality over age.

TABLE 1 ABOUT HERE

The major alternative to the derivatives of mortality, in conceptualizing deceleration, is the slope of the natural log of mortality, dubbed the Lifetable Aging Rate (LAR) by Horiuchi (e.g., 1997). Its chief disadvantage for this paper is that, because the LAR is relative to the overall level of mortality, mortality acceleration/deceleration as measured with the LAR is sensitive to the level of the age-invariant component of mortality (Vaupel and Zhang 2010). In contrast, the mortality derivatives (as will be seen) are functions only of the derivatives of cohort frailty composition and subpopulation mortality. Using the derivatives of mortality in its own scale therefore allows us to focus more cleanly on the contribution of mortality selection.

Three common assumptions about heterogeneity and deceleration

Demographic work on mortality selection has been considerably advanced by efforts to articulate explicit intuitions about the conditions for mortality to decelerate. The purpose of this paper is to show that three common such intuitions are wrong.

First, demographers frequently adopt the heuristic that mortality decelerates when the percent frail in the cohort reaches some low critical value. That heuristic seems to underlie some analyses of deceleration, particularly those that attempt to explicitly relate deceleration patterns to social inequality. For example, Lynch et al. (2003), which significantly advances the

deceleration literature by providing one of the most explicit discussions of how characteristics of population heterogeneity affect deceleration timing, argues that:

> A population with a large number of frail members relative to robust members will experience deceleration when mortality rates are higher (and potentially at a later age) than a population whose membership is equally distributed across frail and robust groups. *In the former case, it simply takes longer for mortality to select out the frailer members.* [Lynch et al. 2003: 462; emphasis added][2]

Heathcote, Puza, and Roberts (2009) makes a similar claim in a paper giving what seems to be the first systematic consideration of the possibility that selection-induced deceleration may be followed by *reacceleration* in populations composed of several subgroups. Their paper shows that:

> [T]here exist models of mixtures of Gompertz groups such that, depending on the extent of heterogeneity, there may be none, one or several age intervals of deceleration of the population hazard function interspersed with intervals of acceleration. Gompertz-like behaviour may then be resumed at extreme old age. *An intuitive explanation is that deceleration occurs when the weakest group is dying out, followed by a brief assertion of*

[2] Similarly, Lynch and Brown (2001) describes the general result that high-mortality populations decelerate at younger ages like this:

> The heterogeneity hypothesis of Horiuchi and Wilmoth (1998) suggests that the age at which deceleration begins should increase over time. The rationale for this prediction is that, as a population becomes more homogeneously robust, the frailer members of the population live longer. Hence their mortality patterns are more similar to that of the most robust subpopulation. *This implies a later age before mortality rates come to be governed by the more robust subpopulation, and hence an older age at which deceleration begins.* [Lynch and Brown 2001: 81; emphasis added]

> *Gompertz acceleration before the next weakest group dies out, and so forth.* [Heathcote et al. 2009: 482; emphasis added]

Second, demographers commonly assume that mortality decelerates only once—or at least, that it can decelerate only once if there are no more than two homogeneous subpopulations. The assumption of a single deceleration is built into the standard parametric form used to model mortality deceleration, the logistic model (e.g., Bongaarts 2005, Kulminski et al. 2007, Rau et al. 2009, Thatcher 1999), as well as the alternative arctangent form used by Lynch and Brown (2001; Lynch et al. 2003). The expectation that multiple decelerations are precluded in a two-subpopulation model is made explicit in Heathcote et al. (2009), which proposes that, in a cohort with k heterogeneous closed subpopulations, mortality may decelerate in a maximum of k-1 intervals.[3]

Finally, it is common to conceptualize the derivatives of aggregate mortality as a competition between subpopulation acceleration, which leads the aggregate hazards to accelerate, and the declining frailty composition (driven by mortality selection) of the population, which leads the aggregate hazards to decelerate. Thus, reacceleration, when it is considered, is assumed to reflect the accelerating mortality of subpopulations, overwhelming the decelerating effect of selection. Mortality selection per se is assumed to produce only deceleration, never acceleration.

In what follows, I show that each of these assumptions can fail.

[3] The results in the present paper do not directly speak to this proposal because Heathcote et al.'s model assigns heterogeneous slopes to the subpopulations, whereas the model presented here assumes proportional hazards. This paper shows that in the proportional hazards setting, generally considered a more restrictive assumption, even cohorts with only two, not three, closed

An example: High-frailty deceleration and multiple deceleration

To make concrete what follows, I begin by introducing as a running example a single simulated cohort, drawn from a class of simulated cohorts described in detail below. This cohort consists of two subpopulations, each with Gompertz mortality, with 75 percent of the cohort in the frail subpopulation at age 50. (I justify the reasonableness of this example, in terms of its parameters and its similarity to real data, when I describe the simulations below. The purpose of this section is simply to fix intuitions about what the phenomena explored in this paper can look like.)

Figure 1 displays this example cohort from ages 50-100 (by which age the frail are virtually extinct, with the aggregate annual hazard equal to .34). The left column gives the frailty composition (proportion frail) of the cohort, and the right column, cohort mortality; the first row presents those quantities over age, while the second, third, and fourth rows give those quantities' respective first, second, and third derivatives over age. In all panels, the dashed dark lines represent Gompertz mortality; the thick grey lines, relative deceleration; and the thick black lines, absolute deceleration. The dashed light vertical line in all panels marks the point where the frail are exactly half of the cohort (we will see that some selection dynamics are importantly different on each side of this line). In the bottom four panels, the dashed light horizontal line marks the point where the second and third derivatives change sign.

FIGURE 1 ABOUT HERE

Panel A of Figure 1 shows the frailty composition of the cohort. While this line is always declining, we will see that the changing speed of its decline drives much of the results to follow. Panel E, in the upper right, which gives the aggregate mortality of the cohort over age, illustrates the counter-intuitive patterns at the heart of this paper. In this cohort, mortality accelerates

---

subpopulations can experience two successive decelerations.

exponentially until age 68, with 66 percent of the cohort frail, when the first interval of relative deceleration begins. At age 75, with 54 percent of the cohort frail, mortality decelerates absolutely; the second derivative remains negative until age 84 (16 percent frail), when a second period of Gompertz mortality begins. This persists until age 91 (nine-tenths of one percent frail), when mortality again decelerates relatively until age 94 (two-tenths of one percent frail), when the cohort enters Gompertz mortality for the third and final time.

The bottom panels of Figure 1, giving the derivatives of frailty composition and cohort mortality, are discussed below when I present the equations for each of these derivatives. Summarizing numerically the key information from Figure 1, Table 2 lists the points, in age and in frailty composition, at which the second and third derivatives of frailty composition and of mortality switch sign. I turn now to the analytical investigation of these derivatives.

TABLE 2 ABOUT HERE

Analytical results:

High-frailty deceleration, multiple deceleration, and selection-driven acceleration are possible in principle

Here I present the first, second, and third derivative over age of cohort mortality. These derivatives are helpful in generating intuitions for how high-frailty deceleration, multiple deceleration, and selection-driven reacceleration are possible even in an exceedingly simply mortality system. Intuitions about deceleration commonly focus on the number of surviving frail members of a cohort, relative to the robust. But, as we will see, it is not directly the relative size

of the frail subpopulation, but rather its rate of decline, that drives deceleration. This rate of decline turns out to have a non-monotonic relationship to the relative size of the frail subpopulation, leading the rate of decline of the cohort's frailty composition—and, sometimes, cohort mortality—to accelerate and decelerate in unexpected ways as the cohort ages. The results given here extend the classic Heterogeneity's Ruses (Vaupel and Yashin 1985), underscoring the dynamism and complexity of the role of mortality selection in both producing and mitigating mortality deceleration.

Consider a mortality model in which frailty is dichotomous and fixed, and in which the resulting two closed subpopulations—the frail and the robust—have proportional Gompertz hazards (one of the simple models in Vaupel and Yashin's [1985] classic mortality selection paper). Equation 1 gives each subpopulation's mortality, $\mu(x)$, at age x:

$$\mu_r(x) = \alpha e^{\beta \cdot x} \tag{1a}$$

$$\mu_f(x) = f\alpha e^{\beta \cdot x} = f \cdot \mu_r, \quad f > 1 \tag{1b}$$

Equation 1a gives the mortality of the robust subpopulation as a function of age, x (which can be defined as years since a baseline age, which may or may not be birth), with intercept α and log-linear slope β. Equation 1b gives the age-specific mortality of the frail subpopulation as the mortality of the robust times a frailty multiplier f, which exceeds 1. Thus, α is the mortality of the robust at the baseline age, and f is the ratio of frail to robust mortality at any age.

Equation 2 gives the mortality of the cohort as a whole as the mortality of each subpopulation, weighted by their share of the cohort:

$$\begin{aligned}\overline{\mu}(x) &= \pi(x)\mu_f(x) + \big(1-\pi(x)\big)\mu_r(x) \\ &= \mu_r\big[1+\big(f-1\big)\pi(x)\big]\end{aligned} \tag{2}$$

where $\pi(x)$ is the *frailty composition*, the proportion of the cohort that is frail. Thus, the derivatives of cohort mortality are functions of the derivatives of the mortality of each subpopulation and of the frailty composition, the latter increasing in importance when the difference between frail and robust mortality is large.

Equation 3 gives the slope of mortality with respect to age:

$$\overline{\mu}'(x) = \pi'(x)\big(\mu_f(x)-\mu_r(x)\big) + \pi(x)\mu_f'(x) + \big(1-\pi(x)\big)\mu_r'(x) \tag{3}$$

Readers may recognize this expression for $\mu'(x)$ as a special case of Vaupel and Zhang's (2010) elegant result that the slope of mortality at any age is the average slope of the two subpopulations ($\pi(x)\mu_f'(x) + \big(1-\pi(x)\big)\mu_r'(x)$) minus the variance of mortality at that age. Here, that negative variance is expressed as the difference in subpopulation mortalities ($\mu_f(x)-\mu_r(x)$) weighted by the slope of frailty composition—i.e. the rate of decline in the percent frail—at that age, $\pi'(x)$. Panel F of Equation 1 displays the slope of mortality for the aggregate cohort; its two turning points will become more interpretable as we analyze the slope of frailty composition.

Equation 4 gives the slope of frailty composition with respect to age:

$$\pi'(x) = -\pi(x)\big(1-\pi(x)\big)\big(\mu_f(x)-\mu_r(x)\big) \tag{4}$$

The absolute value of this expression, which is always negative (i.e., the proportion of the cohort that is frail is always declining, until the frail become extinct), is the *rate of frailty decline*. The

rate of frailty decline can be thought of as the *intensity of mortality selection*. This rate of frailty decline is the quantity on which the results in this paper hinge. Its key characteristic is that the $\pi(x)\left(1-\pi(x)\right)$ term is at its maximum when $\pi(x)=.5$. Thus, all else equal, the selection of the frail out of the cohort is most intense when the frail are half the population, and less intense when the frail are either a large majority or a small minority. However, in a real cohort, all else is *not* equal: the $\mu_f(x)-\mu_r(x)$ term is greatest at high values of subpopulation mortality, which occur at older ages. Thus, over age, the $\mu_f(x)-\mu_r(x)$ term is always increasing, while the $\pi(x)\left(1-\pi(x)\right)$ term (henceforth expressed more compactly, with its negative sign, as $\pi^2(x)-\pi(x)$) increases if the frail are a majority of the population (so that loss of the frail moves the cohort closer to $\pi(x)=.5$) and decreases over age once the frail have become a minority. Thus, in the example cohort, as shown in Panel B of Figure 1, the first derivative of frailty composition reaches its minimum—that is, the rate of frailty decline is at its maximum—at age 82, when 27 percent of the cohort is frail. The rate of frailty decline approaches zero toward both extremes of frailty composition. These simple dynamics drive the surprising complexity investigated in this paper.

Equation 5 gives the second derivative of frailty composition with respect to age:

$$\begin{aligned}\pi''(x) = &\left(\pi^2(x)-\pi(x)\right)\left(\mu_f'(x)-\mu_r'(x)\right)\\ &+\pi'(x)\left(2\pi(x)-1\right)\left(\mu_f(x)-\mu_r(x)\right)\end{aligned} \tag{5}$$

When this expression is negative, the rate of frailty decline is increasing; when positive, the rate of frailty decline is decreasing. The first term of Equation 5, $\left(\pi^2(x)-\pi(x)\right)\left(\mu_f'(x)-\mu_r'(x)\right)$, is always negative. The second term, $\pi'(x)\left(2\pi(x)-1\right)\left(\mu_f(x)-\mu_r(x)\right)$, is negative when $\pi(x)\geq .5$,

but positive when $\pi(x) < .5$, since $2\pi(x) - 1$ switches sign at $\pi(x) = .5$. Thus, the rate of frailty decline can in principle slow down, but only when the frail are a minority of the population. While the frail remain the majority, the rate of frailty decline—that is, the intensity of selection for robustness in the cohort—is always increasing over age. The reflection of the declining rate of frailty decline after age 82, and frailty composition .27, in the positive sign of the second derivative of frailty composition is shown in Panel C of Figure 1 and in the upper left panel of Table 2. Equations 4 and 5 illuminate the complex role of mortality selection, as the increasing and then decreasing rate of frailty decline contributes to mortality deceleration until age 82, and mortality acceleration afterward.

Equation 6 gives the third derivative of frailty composition with respect to age:

$$\begin{aligned}\pi'''(x) = {} & \pi''(x)\left(2\pi(x)-1\right)\left(\mu_f(x)-\mu_r(x)\right) + 2\left(\pi'(x)\right)^2\left(\mu_f(x)-\mu_r(x)\right) \\ & +2\pi'(x)\left(2\pi(x)-1\right)\left(\mu_f'(x)-\mu_r'(x)\right) + \left(\pi^2(x)-\pi(x)\right)\left(\mu_f''(x)-\mu_r''(x)\right)\end{aligned} \tag{6}$$

When this expression is negative, the rate of frailty decline is accelerating; when it is positive, that rate, and hence mortality selection, is decelerating. The second term of Equation 6, $2\left(\pi'(x)\right)^2\left(\mu_f(x)-\mu_r(x)\right)$, is always positive, and the fourth, $\left(\pi^2(x)-\pi(x)\right)\left(\mu_f''(x)-\mu_r''(x)\right)$, is always negative. The third term, $2\pi'(x)\left(2\pi(x)-1\right)\left(\mu_f'(x)-\mu_r'(x)\right)$, switches sign at $\pi(x) = .5$, negative when the frail are a majority of the cohort and positive when they are a minority. The first term, $\pi''(x)\left(2\pi(x)-1\right)\left(\mu_f(x)-\mu_r(x)\right)$, is also always negative when $\pi(x) \geq .5$, but can have either sign when $\pi(x) < .5$. The positive second term—the only term that is positive when $\pi(x) > .5$—reflects that the absolute value of $2\pi(x) - 1$ is smallest when $\pi(x)$ is near .5, and

largest at the extremes (unlike the other terms on the right side of Equation 5, which increase in absolute value as the frailty composition of the cohort declines toward half). In short, on either side of $\pi(x) = .5$, whether the rate of frailty decline is accelerating or decelerating depends on the other parameter values: the intercept of robust mortality, log-slope of robust and frail mortality, and frailty multiplier. In fact, in the example cohort, as shown in Panel D of Figure 1 and the lower left panel of Table 2, the third derivative of frailty composition switches sign just after $\pi(x) = .5$: it is negative until age 77, when the frail are 47 percent of the cohort. It then remains positive until age 87, when only 7 percent of the cohort is frail, after which point in remains negative but approaches zero as the frail become extinct.

Returning to cohort mortality, the bottom four panels of Figure 1 highlight that the dynamics of the second and third derivatives of cohort mortality (shown on the right), whose signs respectively define absolute and relative deceleration, are heavily driven by the second and third derivatives of frailty composition (shown on the left). Equation 7 gives the second derivative of cohort mortality with respect to age:

$$\begin{aligned}\bar{\mu}''(x) = {} & \pi''(x)\left(\mu_f(x) - \mu_r(x)\right) + 2\pi'(x)\left(\mu_f'(x) - \mu_r'(x)\right) \\ & + \pi(x)\mu_f''(x) + \left(1 - \pi(x)\right)\mu_r''(x)\end{aligned} \tag{7}$$

Mortality decelerates absolutely when this expression is negative. The third term of Equation 7, $\pi(x)\mu_f''(x) + \left(1 - \pi(x)\right)\mu_r''(x)$, representing the composition-weighted increase in subpopulation acceleration, is always positive, and the second, $2\pi'(x)\left(\mu_f'(x) - \mu_r'(x)\right)$, representing the difference between the frail and robust subpopulation slopes weighted by twice the rate of frailty

decline, is always negative. The first term, $\pi''(x)\left(\mu_f(x)-\mu_r(x)\right)$, has the sign of the second derivative of frailty composition: it is always negative when the frail are a majority, $\pi(x)\geq .5$, but can be positive or negative when the frail are a minority, $\pi(x)<.5$. In principle, then, mortality can decelerate absolutely when the frail are either a majority or a minority of the cohort. In the example cohort, as shown in Panel G of Figure 1 and in the upper right panel of Table 2, mortality decelerates relatively at age 75, when the frail are 54 percent of the cohort, and reaccelerates at age 84, when the frail are 16 percent of the cohort.

Equation 8 gives the third derivative of cohort mortality with respect to age:

$$\begin{aligned}\bar{\mu}'''(x) = {} & \pi'''(x)\left(\mu_f(x)-\mu_r(x)\right)+3\pi''(x)\left(\mu_f'(x)-\mu_r'(x)\right) \\ & +3\pi'(x)\left(\mu_f''(x)-\mu_r''(x)\right)+\pi(x)\mu_f'''(x)+\left(1-\pi(x)\right)\mu_r'''(x)\end{aligned} \tag{8}$$

Mortality decelerates relatively when this expression is negative. The fourth term, $\pi(x)\mu_f'''(x)+\left(1-\pi(x)\right)\mu_r'''(x)$, representing the composition-weighted increase in subpopulation jerk, is always positive, and the third, $3\pi'(x)\left(\mu_f''(x)-\mu_r''(x)\right)$, representing the difference between frail and robust acceleration weighted by three times the rate of frailty decline, is always negative. Both the first and second term are always negative when the frail are a majority, $\pi(x)\geq .5$, and may take either sign when the frail are a minority, depending respectively on the signs of the third and second derivatives of frailty composition. In the example cohort, as shown in Panel H of Figure 1 and in the lower right panel of Table 2, mortality decelerates relatively at age 68, when the frail are 66 percent of the cohort; reaccelerates at age 81, when the frail are 31 percent; decelerates relatively a second time at age 91, when the frail are only 1 percent; and

reaccelerates a final time at age 94, when the frail are only two-tenths of one percent of the cohort.

These equations generate some intuition for how mortality may decelerate while a majority of the cohort is frail—as the rate of frailty decline increases, with frailty composition hurtling downward toward half of the cohort—and evince a complex pattern of acceleration and deceleration when the frail are a minority of the cohort. Yet the relative complexity of the expressions suggests that the ultimate patterns may depend heavily on the values of the subpopulation mortality parameters. Therefore, to more systematically assess what we should expect in human cohorts, I turn to simulations.

Simulation results:

High-frailty deceleration, multiple deceleration, and selection-driven acceleration occur widely in model populations compatible with the Human Mortality Database

Four parameters define the mortality model: the intercept for robust mortality, α; the log-slope of mortality for both subpopulations, β; the frailty multiplier, or ratio of frail to robust mortality at any age, f; and the baseline percent frail, π_0. All four parameters are unobserved in real data, since subpopulation membership is latent by assumption. Thus, in generating realistic simulations, the goal is to find parameter combinations that generate aggregate cohorts whose parameters match those of real human cohorts (and, perhaps, whose subpopulation parameters fit some theoretically-driven idea of what is reasonable). It is easiest to take a brute force approach

to this problem, by generating many latent subpopulation models and keeping only those whose aggregate parameters are consistent with the aggregate life tables of known human cohorts.

To limit the complexity over four dimensions, one parameter—baseline frailty composition—is restricted to constant values in the main simulations. The baseline frailty composition is set across all simulations to .75, a high value chosen to make visible the selection dynamics when frailty is common as well as rare. The baseline age is 50, which leaves the model agnostic as to whether mortality rises during late adulthood with the same log-slope as it had earlier in life.[4] The model therefore assumes that 75 percent of the population surviving to age 50 is frail.[5] Thus, these simulations represent cohorts in which mortality advantage, rather than disadvantage, is the exceptional condition. Such populations are easily imagined; for example, Lynch et al.'s (2003) study of African-Americans born 1870-1972 hypothesizes that this population, due to its extreme deprivation, was nearly homogeneously frail.[6] The frailty

[4] The mortality derivatives are evaluated up to age 150, by which point the frail are extinct in all cohorts, to ensure that no periods of deceleration or reacceleration are censored. However, parametric (Gompertz and, later, logistic) models used to assess the similarity between real and simulated data are estimated on ages 50-100 to ensure comparability between the models for real and simulated cohorts (since real data do not extend to age 150).

[5] One might be concerned that it is impossible for a cohort to be 75 percent frail at age 50 with reasonably-valued Gompertz subpopulations because too many frail will have died by age 50. It turns out that this is not the case. Were the subpopulation intercept and slope parameters constant from birth, this would correspond to a proportion frail at birth in the range of .750 to 1 in the final universe of simulated cohorts, with a mean value of .887 (calculations omitted; available upon request). A proportion frail of 1 is incompatible with the assumption of two subpopulations. The 37 cohorts that generate that result, given the assumption of constant lifetime subpopulation mortality parameters, are the cohorts with the lowest β (slope) and highest α (intercept) values in the simulation universe. Excluding them does not appreciably change results.

[6] Another example might be mortality data that excludes certain dimensions of extreme social stratification. For example, mortality data from apartheid-era South Africa, if not stratified by race, could be conceived of as an aggregation of a large, high-mortality Black subpopulation and a smaller, advantaged White subpopulation, as well as an intermediate Coloured subpopulation. Since frail and robust are relative categories, relevant examples are ones in which the best dichotomization of mortality risk puts most of the population into the higher-mortality group, but

multiplier is modeled at eight values, ranging in units of .5 from 1.5 to 5; the low end represents fairly modest disadvantage, while the top end is at the extreme of what we might consider plausible for human populations.[7] The intercept for the robust and the log-slope for the two subpopulations are varied nearly continuously, in increments of .001—in the range [.001,.2] for α and [.001,.4] for β. In total, this produces 640,000 simulated cohorts before evaluating the resulting parameters for plausibility.

To winnow these 640,000 simulated cohorts down to a realistic subset, I estimate a Gompertz model on each aggregated cohort,[8] and keep only the ones that fall inside a convex hull formed around the intercept and log-slope parameters estimated from the 2,352 historical European cohorts collected in the Human Mortality Database (HMD).[9] Figure 2, Panel A, displays in grey the aggregate Gompertz parameters of the resulting universe of 1,151 simulated

---

not necessarily ones in which most people are "frail" in some absolute sense.

[7] For a rough-and-ready sense of what fairly extreme mortality differentiation looks like, consider sex differences in Russian mortality. Russian cohorts born 1872-1980 have an age-specific ratio of the male to female annual mortality rate ranging between .77 to 4.87, with a mean (weighted by total exposure) of 2.85. The sex ratio is increasing over time; for cohorts born beginning in 1950, the mean weighted ratio is 3.01, and when limited to ages 50-100, as in the simulations, the ratio for those modern cohorts is 3.19 (author's calculations from Human Mortality Database data).

[8] Estimating Gompertz models (and, later, logistic models) on the simulated cohorts requires estimating discrete survivorship at each age so that the parametric estimation can be weighted by survivorship, as in real data on individuals. These discrete survivorships are estimated from the mortality functions using standard life table methods that assume constant mortality within each age interval (Preston et al. 2001: 46-47). To make palatable this assumption, which violates the assumption of Gompertz subpopulation mortality, I use age increments of only four days.

[9] Human Mortality Database. University of California, Berkeley (USA), and Max Planck Institute for Demographic Research (Germany). Available at http://www.mortality.org or http://www.humanmortality.de (data downloaded on August 18, 2011). The data represent all cohort (vs. period) data included in the HMD.

cohorts that meet these criteria, with those of the HMD cohorts overlaid in black.[10] All simulation results are drawn from this universe.

FIGURE 2 ABOUT HERE

Returning now to the example cohort presented in Figure 1 and Table 2: that cohort, defined by the parameters f=5, α=.002, and β=.103, has Gompertz intercept .009 and slope .081. It was chosen arbitrarily from among those exhibiting multiple relative decelerations whose aggregate parameters fell in a dense cluster of HMD cohorts, born in Sweden and Denmark in the mid-19th Century and in England, Wales, and Scotland in the late 19th Century.

Simulation results

Relative and absolute deceleration are both rampant in this universe, and whether they occur is closely predicted by the frailty multiplier, as summarized in the first two columns of Table 3. None of the simulated cohorts in this universe with frailty multiplier f=1.5 decelerate. In contrast, as shown in the first column of Table 3, all of the cohorts in which the frail subpopulation has at least three times the mortality of the robust, $f \geq 3$, and some in which the frail have only two and a half times the mortality of the robust, f=2.5, decelerate absolutely. These absolute decelerations occur in the age range 68 to 90, at annual mortality values ranging from .07 to .16. As shown in the second column of Table 3, *all* of the 910 cohorts in which the frail subpopulation has at least

[10] The historical sweep of the HMD cohorts is from the lower right (high intercept, low Gompertz slope) to the upper left (low intercept, high slope). The increasing slope over time presumably reflects diminished mortality selection in childhood, so that a greater proportion of relatively frail cohort members survive to old age, contributing to mortality compression (Kannisto 2000).

twice the mortality of the robust, $f \geq 2$, exhibit at least one relative deceleration. These relative decelerations occur at ages between 61 and 105, at mortality values ranging from .04 to .20. Mortality decline is not found in the universe of simulated cohorts close to the HMD cohorts.

TABLE 3 ABOUT HERE

High-frailty and multiple deceleration—Panels B and C of Figure 2 shows the incidence of high-frailty deceleration, which I define as deceleration when the frail are a majority of the cohort. Panel B shows simulated cohorts evincing high-frailty absolute deceleration. Such decelerations are found in 30 percent of total cohorts in this universe, and 56 percent of those with any absolute deceleration; in other words, most absolute decelerations occur while the frail are a majority of the cohort. High-frailty absolute decelerations occur at ages ranging between 68 and 85. These high-frailty absolute decelerations are found in cohorts with frailty multipliers ranging from 3.5 to 5—indeed, as shown in the third column of Table 3, *all* cohorts with frailty multiplier equal to at least 4 evince high-frailty absolute deceleration.

Panel C of Figure 2 shows that high-frailty relative deceleration can occur across the full range of intercept and slope values derived from the HMD cohorts. It occurs here at ages ranging from 61 to 90. Seventy-nine percent of total cohorts in this universe evince high-frailty relative deceleration. More strikingly, as shown in the fourth column of Table 3, 100 percent of cohorts with any relative deceleration at all—that is, all and only cohorts with a frailty multiplier of at least 2—decelerate relatively while most of the cohort is frail.

Panel D of Figure 2 shows the incidence of multiple deceleration, which occurs only for relative, not absolute, deceleration. Multiple relative decelerations occur in cohorts at many points across the range of intercept and slope parameters, though much more sparsely than high-

frailty decelerations. As shown in the fifth column of Table 3, all and only cohorts with frailty multiplier equal to 4.5 or 5—fifteen percent of the universe with aggregate parameters similar to HMD cohorts—have two intervals of relative deceleration, one when the frail are a majority of the cohort, and a second one when they are a small minority.

At what percent frail does deceleration and reacceleration occur in this universe? Figure 3 displays the frailty composition at absolute deceleration, relative deceleration as a whole, and relative deceleration restricted to cohorts that decelerate only once (each with their respective reaccelerations reflected in the bottom row). The results underscore the problems for the heuristic that mortality decelerates only when the frail are nearly depleted. Panel A of Figure 3 shows that absolute deceleration *never* corresponds, in these simulations, to that heuristic: absolute deceleration can occur when the frail are a majority or a minority, but never occurs here when they are less than 35 percent of the population. As shown in Panel B, reaccelerations following absolute deceleration, likewise, occur well before frailty depletion, when the frail are between 15 and 34 percent of the cohort. Panels C and D, which show all relative decelerations, suggest that relative deceleration does sometimes occur as the heuristic would predict, with the cohort decelerating when the frail are nearly gone (in these cohorts, when frailty composition ranges between half a percent and one and two-tenths percent) and reaccelerating shortly thereafter. But Panels E and F demonstrate that this pattern occurs only among the second of two decelerations. These panels are limited to the 67 percent of simulated cohorts with only one relative deceleration, that is, all and only cohorts where the frailty multiplier is between 2 and 4 (inclusive). They show that when there is a single relative deceleration, the deceleration occurs when most of the cohort is frail. In short, neither absolute nor relative deceleration corresponds to the conventional picture of a single deceleration when the frail are approaching extinction.

FIGURE 3 ABOUT HERE

Selection-driven acceleration—It seems intuitive to conceive of deceleration/acceleration as a tradeoff between declining frailty composition (produces deceleration) and accelerating subpopulations (produces acceleration). Yet a consequence of the results given above is that—because deceleration does not in general occur when the frail are nearly extinct, and because the slope of frailty decline is smaller at the extremes of frailty composition—declining frailty composition can also produce *acceleration*.

To underscore this point, I offer an artificial calculation as a device for isolating the role of declining frailty composition, illustrating what I call *selection-driven acceleration*. Starting from the example cohort given above (π_0=.75, f=5, α=.002, and β=.103), imagine that we could hold subpopulation mortality and its derivatives fixed at their levels in age 81—the age when the aggregate second derivative reaches its minimum (with 31 percent of the cohort frail)—while the percent frail is left varying as in the actual cohort. This calculation, which of course reflects a physical and conceptual impossibility, isolates the effects of the declining percent frail from those of increasing subpopulation hazards, slopes and accelerations. Figure 4 shows the results: while mortality in this exercise does not reaccelerate nearly as dramatically without the changes in the subpopulation mortality of the real cohort, nevertheless it does reaccelerate. As we would expect from Equation 4, the rate of frailty decline at all ages (in this post-deceleration age range beginning at age 81) is smaller for the artificial simulation than in the real model cohort with changing subpopulation mortality, since in that real model, the increasing difference over age between frail and robust mortality increases the rate of frailty decline. Most importantly, in the artificial simulation, the rate of frailty decline is decreasing (the second derivative of frailty composition, given in Equation 5, is positive beginning at age 85) as frailty composition falls

farther below .5. This leads mortality to reaccelerate even in the absence of subpopulation-level changes. The reacceleration is generated entirely by the declining percent frail, that is, by mortality selection. This demonstrates that mortality selection can contribute to acceleration as well as deceleration.

FIGURE 4 ABOUT HERE

Further implications

Problems of estimating age at deceleration with common parametric models

Results so far suggest that, even in an exceedingly simple model, it is possible, indeed plausible, for populations to experience at least one period of deceleration followed by reaccelerating mortality. Yet the standard parametric forms used to model older ages—most often, logistic models (e.g., Bongaarts 2005, Kulminski et al. 2007, Rau et al. 2009, Thatcher 1999), and occasionally, very similar arctangent models (Lynch and Brown 2001; Lynch et al. 2003)—assume that mortality decelerates at most once and never reaccelerates. It turns out that this can lead such parametric forms to systematically misestimate deceleration timing.

Figure 5 plots the deceleration timing derived from logistic models against the actual timing of the simulated cohorts, with the main diagonal provided as a reference line.[11] Panel A shows the results for absolute deceleration. The results indicate that the logistic models badly overestimate the age of deceleration for these cohorts. The degree of overestimation ranges

[11] Deceleration timing for logistic models is defined in the same way as for the nonparametric measures—that is, when the second or third derivative become negative—using formulas for

between 18 and 39 years (28 years on average). Panels B, C, and D show the results for relative deceleration, considering, respectively, single, first, and second relative decelerations. Thus, panel B gives actual vs. estimated deceleration age for those cohorts that decelerate only once. Panels C and D each give the same outcome measure—deceleration timing estimated from logistic models on cohorts that decelerate twice—spread over two different regions on the horizontal axis, reflecting the cohort's two deceleration points. The graphs show that the logistic models fit poorly the deceleration patterns for first and single decelerations, and fit well second decelerations. Single deceleration ages are overestimated by 5 to 27 years (16 on average), and first decelerations by 17 to 28 years (22 on average). In contrast, the logistic models *under*estimate the age at second relative decelerations by between 1 and 6 years (underestimating by 3 on average), a far smaller difference. The extreme overestimation of the age at absolute deceleration partly follows from this, since absolute deceleration, when it occurs, follows relative deceleration, but in these cohorts absolute deceleration follows only the high-frailty relative deceleration, whereas logistic models far more closely track the timing of the (relatively rare) low-frailty relative deceleration.

FIGURE 5 ABOUT HERE

Most troublingly, perhaps, the logistic models falsely detect deceleration with alarming frequency. Absolute deceleration is detected in all cohorts that decelerate only relatively. Most strikingly, relative and absolute deceleration are predicted in all cohorts that do not decelerate at all (that is, in this simulation universe, all and only cohorts with frailty multiplier f=1.5). The

---

those derivatives taken from Rau et al. (2009).

estimated age of relative deceleration for these cohorts ranges from 83 to 100 (mean 90), and absolute deceleration from 99 to 117 (mean 107).

These falsely detected decelerations are especially problematic because the ages at which deceleration is identified even though it does not occur are similar to the ages at which previous research using logistic and similar parametric models has identified deceleration. Using logistic models, Rau et al. (2009) find relative deceleration among English and Welsh women at age 93, and absolute deceleration at age 103, and Bebbington et al. (2007) find absolute deceleration among Canadian men at age 92 and women at age 96.5. It stands to reason that similar problems might occur using arctangent models due to their similarity to logistic models (Lynch and Brown 2001); using arctangent models, Lynch and Brown (2001) find absolute deceleration among white women in the U.S. at ages ranging 95-96 and white men at ages 93-95, from 1968-1992; and Lynch et al. (2003) find absolute deceleration for U.S. whites at ages ranging 93-95 and for blacks at ages 92-96 (in unadjusted data) or 101-104 (in data adjusted for potential misreporting), from 1970-1992.

Moreover, these previous results highlight that studies using differences in deceleration timing to understand heterogeneity between and within populations are often based on differences in deceleration onset (between groups or over time) of only a few years, much smaller than the error in most estimated deceleration timing for the simulated cohorts in this paper. These results collectively suggest that demographers whose primary object of study is deceleration should not rely on single-peak parametric models, and instead should strive to substitute conventional parametric approaches with much more flexible, ideally nonparametric, models.

Comparing deceleration across cohorts

A central motivation for accurately measuring deceleration timing is that comparing deceleration timing across cohorts—whether distinct birth cohorts or the same birth cohort across multiple closed social groups—may permit demographers to infer something about the differences in the cohorts' distributions of mortality risk. This important endeavor (pursued in Horiuchi and Wilmoth 1997, 1998; Lynch and Brown 2001; Lynch et al. 2003) links the measurement of deceleration to the study of inequality and change in mortality.

This approach necessitates a qualitatively simple relationship between unobserved heterogeneity patterns and observed patterns of deceleration timing. Accordingly, such reasoning was advanced considerably by Lynch et al. (2003), which articulated explicit predictions about the circumstances in which one cohort should decelerate at an older age and higher mortality level than another, and used those predictions to infer changes in mortality heterogeneity within racial groups from changing patterns of deceleration.[12] One of these predictions was quoted above: "A population with a large number of frail members relative to robust members will experience deceleration when mortality rates are higher (and potentially at a later age) than a population whose membership is equally distributed across frail and robust groups" (Lynch et al. 2003: 462). Earlier I suggested that this prediction was grounded in the heuristic that mortality decelerates only when the frail are largely extinct. In this section, I show more explicitly how this prediction will sometimes fail.

[12] Horiuchi and Wilmoth [1998] makes a similar contribution for deceleration patterns across causes of death.

First, for concreteness, consider an example. I hold fixed the subpopulation parameters at frailty multiplier f=5, robust intercept α=.004, and log-slope β=.09 and allow the baseline frailty composition, π_0, to vary in units of .05 from .05 to .75.[13] Figure 6 displays the results. The rows show, respectively, the mortality, age, and percent frail of the cohort at deceleration, plotted against the baseline percent frail. The type of deceleration varies across columns: absolute deceleration; first and single relative decelerations; second relative decelerations; and all relative decelerations together. The prediction articulated in Lynch et al. (2003) is that the lines in the first row and, more tentatively, the second row should be monotonically increasing: the mortality and, more tentatively, age at deceleration (vertical axes) should increase as the baseline frailty composition (horizontal axis) increases.

FIGURE 6 ABOUT HERE

Four important results are suggested in Figure 6. First, since the lines are not all monotonically increasing, it appears that neither age nor mortality always conforms to the prediction that a cohort with larger baseline frailty composition will decelerate at a later age, with higher mortality. This is the case even when second relative decelerations are considered separately from first relative decelerations. Second, the prediction holds up much better for mortality than for age, since the mortality but not the age at absolute and first relative decelerations is monotonically increasing with the baseline frailty composition. This is in line with the greater confidence Lynch et al. express in the prediction for mortality. Third, a different quantity, the percent frail at the onset of deceleration, does appear to rise monotonically with the percent frail at baseline for each type of deceleration, as long as first and second relative

[13] The parameter values are chosen arbitrarily from among those that generate multiple relative decelerations at more than one value of baseline frailty composition. This is necessary for comparing the timing of second relative decelerations across baseline frailty values.

decelerations are distinguished from one another.[14] This, again, contradicts the reasoning that deceleration timing is determined by how long (in age or in accumulated mortality) it takes for the percent frail to fall to extremely low levels. Fourth, the last column underscores the inferential difficulties posed by the possibility of multiple deceleration. Analysts comparing relative decelerations will not in general know whether both are high-frailty, both are low-frailty, or one is each; the possibility that a cohort with high baseline frailty composition will decelerate at both very high and very low frailty makes it more difficult to assess, by measuring a single deceleration point, what the baseline frailty might have been.

Table 4 extends to all parameter combinations in the simulation universe the test of the relationship between baseline frailty composition and the mortality, age, and frailty composition at deceleration. The cells report the proportion of pairs of cohorts for which (across columns) the mortality, age, or percent frail at deceleration is greater than in the cohort with greater baseline percent frail. For absolute decelerations and single relative decelerations, the prediction fares well (albeit imperfectly) for mortality, and poorly for age: higher baseline frailty is associated with higher mortality at deceleration most of the time, and higher age at deceleration rarely. For both mortality and age, the prediction is consistently validated for first relative decelerations (when these are separated from both second and single decelerations), but consistently disproved for second (low-frailty) relative decelerations. The latter is particularly important since, as shown above, it is these second relative decelerations, where the prediction fares worst, that are most closely matched by the estimated deceleration in logistic mortality models. In contrast to these results for mortality and age, all four of these types of deceleration—absolute, single relative,

[14] Indeed, the relationship between percent frail at baseline and percent frail at deceleration is surprisingly linear: linear regressions of the latter on the former for absolute deceleration and single decelerations (the two types with more than two points) yield R-squared values in excess

first relative, and second relative—occur at a higher percent frail among cohorts that began with a larger percent frail at baseline. Only when first and second relative decelerations are considered together is the relationship broken between percent frail at baseline and percent frail at deceleration.

TABLE 4 ABOUT HERE

The striking relationship between the percent frail at baseline and percent frail at deceleration suggests that further investigation into these dynamics is warranted. Unfortunately, the percent frail at deceleration, unlike aggregate mortality and age, is not observable in real data. Thus, the results for percent frail do not directly aid the project of using observed deceleration patterns to test theories about cohorts' unobserved heterogeneity at baseline, such as the theory investigated by Lynch et al. (2003) that African-American cohorts became less homogeneously frail after the Civil Rights Movement, when improved social and political circumstances may have less sharply curtailed their potential longevity. In fact, the real situation of deceleration analysts is more complicated than Table 4 suggests because real data are truncated at the oldest ages. Table 4 adopts the perspective of an observer who is omniscient as to when deceleration does or does not occur. But demographers using real datasets are never sure whether cohorts that appear not to decelerate in fact decelerate at older ages than those observed in the data. Since decelerations usually occur at relatively high frailty composition, cohorts with low baseline frailty may not decelerate at all, yet may be mistaken for cohorts that decelerate at very late ages. Using traditional reasoning, those cohorts would then be presumed to have had unusually high baseline frailty composition, when the reverse would be true.

---

of 99 percent.

In short, it is not necessarily the case that ordering decelerations across cohorts—by mortality or by age—can reveal which cohort had more frail members at birth, even assuming that the cohorts otherwise share the same mortality parameters. Analysts making such inferences may need to develop more precise hypotheses about the unobserved heterogeneity in the cohorts they study in order to know what to infer from the order of deceleration.

Conclusion

This paper has demonstrated three unexpected facts about mortality deceleration, which together have three broader implications for demographers. It has shown that even within a single cohort composed of just two subpopulations with proportional Gompertz hazards:

1. Mortality can decelerate even while a majority of the cohort is frail (*high-frailty deceleration*).

2. Mortality can then reaccelerate while the frail remain a non-negligible part of the cohort. This occurs because the rate of selection is greatest when half the population is frail, so that—counter-intuitively—selection of the frail out of the cohort can cause acceleration, not only deceleration, as the frailty composition dips below half (*selection-driven acceleration*).

3. Mortality can then decelerate a second time as the frailty composition dips further below half (*multiple deceleration*), before finally reaccelerating as the robust become such a large part of the cohort that their acceleration dominates over the negligible selection that remains possible.

These facts have three important implications. First, the first two facts challenge a conventional heuristic that has anchored important intuitions in previous demographic work, namely: the heuristic that mortality decelerates only as the frailty composition is 'nearly' depleted, and reaccelerates only as the frailty composition is 'nearly all' depleted. This paper shows that, while this pattern does occur, it is not the only—or even the main—pattern of deceleration and reacceleration possible in this simple mortality setting. The contribution of diminishing frailty composition to deceleration and reacceleration is more complex than has been previously articulated in the literature.

Second, the second and third facts suggest that conventional parameterizations of old-age mortality may lead analysts astray. Parametric forms used to identify the timing of deceleration, such as logistic (Bongaarts 2005, Kulminski et al. 2007, Rau et al. 2009, Thatcher 1999) or arctangent (Lynch and Brown 2001, Lynch et al. 2003) forms, assume a single mortality plateau. Not only will such parametric forms miss reacceleration and multiple deceleration when they occur; when such patterns occur, these parametric forms may misstate the timing of any deceleration point, as they average observations whose derivatives are significantly more complex than the forms assume. This is particularly problematic for purposes that compare deceleration timing across cohorts—and such comparisons are the central way that deceleration timing bears on inequality, within and across cohorts.

Finally, the facts together qualify the link between deceleration patterns and inequality in one additional way. It turns out that—in contrast to an earlier prediction (Lynch et al. 2003) used to link deceleration patterns to changing heterogeneity among blacks and whites in the United States—all else equal, a cohort with greater baseline frailty composition can decelerate at *lower mortality* and *younger ages* than one with fewer frail members at baseline.

Most of the results presented here apply to a particular kind of population, one in which mortality advantage—longevity relative to one's cohort—is the exception rather than the rule. Some demographic theory on mortality compression suggests that these are likely to be disadvantaged populations, insofar as modern health advances have more dramatically altered mortality by raising much of the population to a higher standard length of life than by allowing the most advantaged to live ever longer (e.g., Brown et al. forthcoming). Thus, it may be among relatively disadvantaged populations that multiple deceleration and high-frailty deceleration may occur. Insofar as such populations often are the least well documented empirically, it may be especially difficult to amass the data required to circumvent the parametric assumptions shown here to sometimes be deeply distorting. On the other hand, recent work examining cross-period and cross-cohort mortality variation at a variety of ages shows that, while mortality advances reduce variation from birth, such advances may increase variation at older ages, in part because with reduced early-life mortality, more frail cohort members live to old age (Engelman, Canudas-Romo and Agree 2010). Thus, even if advantaged populations have fewer frail members from birth than disadvantaged populations, they may have as many or more frail members at the elderly ages in which deceleration may occur. In short, demographic theory does not preclude models with high frailty composition at early-old ages, such as the models explored in this paper, for either disadvantaged or advantaged populations. Moreover, population scientists often wish to compare the mortality of more advantaged and less advantaged groups to one another, and cross-national analyses show that even populations with similar life expectancy may differ considerable in their degree of heterogeneity (Edwards and Tuljapurkar 2005). In practice, then, many comparative analyses will compare the mortality of cohorts that differ in their frailty distribution at whatever age is taken as baseline. The results here suggest that

deceleration may occur at several different points in the process of shifting from a relatively frail to an almost entirely robust surviving cohort. When the cohorts also had very different heterogeneity distributions to begin with, inferences from deceleration patterns to the patterns of heterogeneity within each cohort may be particularly problematic.

What is perhaps most startling is not only that such counter-intuitive patterns are possible, but that they are possible even in an exceedingly simple mortality model. Reality is bound to be more complex, and more complicated models may or may not create even less predictable deceleration dynamics.[15] These results call for caution in modeling and interpreting mortality deceleration, long taken to be a crucial gauge of heterogeneity within cohorts. More broadly, the results in this paper should urge demographers to deepen our theoretical understanding of the surprisingly complex ways that patterns of acceleration and deceleration arise from changing cohort composition. These results highlight the dangers of relying on intuitions about deceleration, since the dynamics of mortality derivatives and mortality selection turn out to be

[15] It is unclear *a priori* how different model assumptions would affect the predictability of deceleration dynamics. We have seen that simple models of binary frailty, such as the present one, readily lead to hard-to-predict multiple deceleration scenarios that can confound common intuitions. I suspect that multivalued discrete frailty scenarios, which may reasonably characterize populations beset by stark inequalities, might similarly produce counterintuitive deceleration results. Smooth frailty assumptions, on the other hand, may simplify matters, as when the computationally convenient gamma-Gompertz model aggregates to a logistic, and hence single-deceleration, mortality regime (Beard 1959, 1971). Gamma-Gompertz models have been widely used in practice (e.g., Gampe 2010, Horiuchi and Wilmoth 1998) ever since Vaupel, Manton, and Stallard (1979) pointed out their convenient properties. The results given here therefore suggest that the choice of an appropriate functional form for unobserved heterogeneity should receive further attention (Steinsaltz and Evans 2004 argue for this position as well). In particular, demographers should investigate options for modeling deeply stratified populations where major dimensions of heterogeneity are unobserved. To the best of my knowledge, all previous models either integrate continuous individual variation at the cost of imposing a unimodal distribution (as in gamma-Gompertz models), or capture the clumping of individual variation that may result from categorical inequalities, such as racial inequality, at the expense of continuous individual variation (in discrete models, such as the model used here).

more complex than anticipated. They suggest a greater need for formal modeling of deceleration dynamics, and in particular, explicitly comparative modeling that matches the kinds of inferences about heterogeneity and inequality for which deceleration patterns have been used as evidence.

REFERENCES

Beard, R. E. 1959. Notes on some mathematical mortality models. *In:* Wolstenholme, G. E. W. & O'Connor, M. (eds.) *The Lifespan of Animals.* Boston: Little, Brown.

Beard, R. E. 1971. Some aspects of theories of mortality, cause of death analysis, forecasting and stochastic processes. *In:* Brass, W. (ed.) *Biological Aspects of Demography.* London: Taylor & Francis.

Bebbington, M., Lai, C.-D., Zitikis, R. 2007. Modeling human mortality using mixtures of bathtub shaped failure distributions. *Journal of Theoretical Biology* 245(3): 528-538.

Bongaarts, J. 2005. Long-range trends in adult mortality: Models and projection methods. *Demography* 42(1): 23-49.

Brown, D. C., Hayward, M., Montez, J. K., Hidajat, M. M., Hummer, R. A., & Chiu, C.-T. N.D.. "The Significance of Education for Mortality Compression in the United States." Unpublished manuscript.

Carey, J. R., Liedo, P., Orozco, D. & Vaupel, J. W. 1992. Slowing of mortality rates at older ages in large Medfly cohorts. *Science,* 258**,** 457-461.

Edwards, R. & Tuljapurkar, S. 2005. "Inequality in lifespans and a new perspective on mortality convergence across industrialized countries." *Population and Development Review*. 31(4): 645-674.

Engelman, M., Canudas-Romo, V., & Agree, E. 2010. "The implications of increased survivorship for mortality variation in aging populations." *Population and Development Review*. 36(3): 511-539.

Fukui, H.H., L. Xiu, & Curtsinger, J.W. 1993. "Slowing of age-specific mortality rates in Drosophila melanogaster." *Experimental Gerontology.* 28(6): 585-99.

Gampe, J. 2010. "Human mortality beyond age 110." In Maier, H., Gampe, J., Jeune, B., Robine, J.-M., & Vaupel, J. (eds). *Supercentenarians (Demographic Research Monographs 7)*. Berlin: Springer. 219-230.

Gavrilov, L. A. & Gavrilova, N. S. 1991. *The biology of life span : a quantitative approach,* New York, Harwood Academic Publishers

Gompertz, B. 1825. On the nature of the function expressive of the law of human mortality, and on a new mode of determining the value of life contingencies. *Philosophical transactions of the Royal Society of London,* 115**,** 513–583.

Heathcote, C. R., Puza, B. D. & Roberts, S. P. 2009. The use of aggregate data to estimate Gompertz-type old-age mortality in heterogeneous populations. *Australian & New Zealand Journal of Statistics,* 51**,** 481-497.

Horiuchi, S. 1997. Postmenopausal acceleration of age-related mortality increase. *Journal of Gerontology: Biological Sciences* 52A: B78-B92.

Horiuchi, S. & Wilmoth, J. R. 1997. Age patterns of the life table aging rate for major causes of death in Japan, 1951-1990. *Journals of Gerontology Series a-Biological Sciences and Medical Sciences,* 52**,** B67-B77.

Horiuchi, S. & Wilmoth, J. R. 1998. Deceleration in the age pattern of mortality at older ages. *Demography,* 35**,** 391-412.

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

Kulminski, A. M., Ukraintseva, S. V., Akushevich, I. V., Arbeev, K. G., & Yashin, A. I. 2007. Cumulative health deficiencies as a characteristic of long life. *Journal of the American Geriatrics Society* 55 (6): 935-940.

Lynch, S. M. & Brown, J. S. 2001. Reconsidering mortality compression and deceleration: An alternative model of mortality rates. *Demography,* 38**,** 79-95.

Lynch, S. M., Brown, J. S. & Harmsen, K. G. 2003. Black-White differences in mortality compression and deceleration and the mortality crossover reconsidered. *Research on Aging,* 25**,** 456-483.

Olshansky, S. J. 1998. On the biodemography of aging: A review essay. *Population and Development Review,* 24**,** 381-393.

Preston, S. H., Heuveline, P. & Guillot, M. 2001. *Demography: Measuring and modeling population processes*. New York: Wiley-Blackwell.

Rau, R., Muszynska, M. & Baudisch, A. 2009. At what age does mortality start to decelerate? Population Association of America. April. Detroit.

Steinsaltz, D., & Evans, S. N. 2004. Markov mortality models: Implications of quasistationarity and initial distributions. Theoretical Population Biology, 65(4): 319-337.

Thatcher, A. R., Kannisto, V. & Vaupel, J. W. 1998. *The force of mortality at ages 80 to 120.* Odense Denmark, Odense University Press

Vaupel, J. W. 1997. Trajectories of mortality at advanced ages. *In:* Wachter, K. W. & Finch, C. E. (eds.) *Between Zeus and the Salmon: The Biodemography of Longevity.* Washington, D.C.: National Academy Press.

Vaupel, J. W. & Gampe, J. 2009. Supercentenarians and the theory of heterogeneity. International Population Conference. September. Marrakech, Morocco.

Vaupel, J. W., Manton, K. G. & Stallard, E. 1979. Impact of heterogeneity in individual frailty on the dynamics of mortality. *Demography,* 16**,** 439-454.

Vaupel, J. W. & Yashin, A. I. 1985. Heterogeneity's ruses: Some surprising effects of selection on population dynamics. *American Statistician,* 39**,** 176-185.

Vaupel, J. W. & Zhang, Z. 2010. Attrition in heterogeneous cohorts. *Demographic Research* 23:26, 737-748.

TABLES

Table 1. Three definitions of deceleration relative to Gompertz mortality

	Mortality, μ(a)		
	Slope, μ'(a)	Acceleration, μ''(a)	Jerk, μ'''(a)
Gompertz mortality	>0	>0	>0
Relative deceleration	>0	>0	**<0**
Absolute deceleration	>0	**<0**	
Mortality decline	**<0**		

Table 2. Sign changes in second and third derivatives of mortality and frailty composition for example cohort

	Mortality			Frailty Composition		
	Age	Frailty	Sign becomes	Age	Frailty	Sign becomes
Second derivative	75	.54	-			
	84	.16	+	82	.27	+
Third derivative	68	.66	-			
	81	.31	+	77	.47	+
	91	.01	-	87	.07	-
	94	.002	+			

Table 3. Deceleration patterns across simulated cohorts, by frailty multiplier

Frailty Multiplier	Absolute deceleration	Relative deceleration	Majority-frail absolute deceleration	Majority-frail relative deceleration	Multiple relative decelerations
1.5	None	None	None	None	None
2	None	All	None	All	None
2.5	Some	All	None	All	None
3	All	All	None	All	None
3.5	All	All	Some	All	None
4	All	All	All	All	None
4.5	All	All	All	All	All
5	All	All	All	All	All

Table 4. Proportion of cohorts for which mortality, age, and percent frail at deceleration increase with baseline percent frail

Deceleration Type	Mortality	Age	Percent Frail
Absolute	.97	.23	1
Relative (single)	.80	.32	1
Relative (first)	1	1	1
Relative (second)	0	0	1
Relative (all)	.80	.35	.96

FIGURES

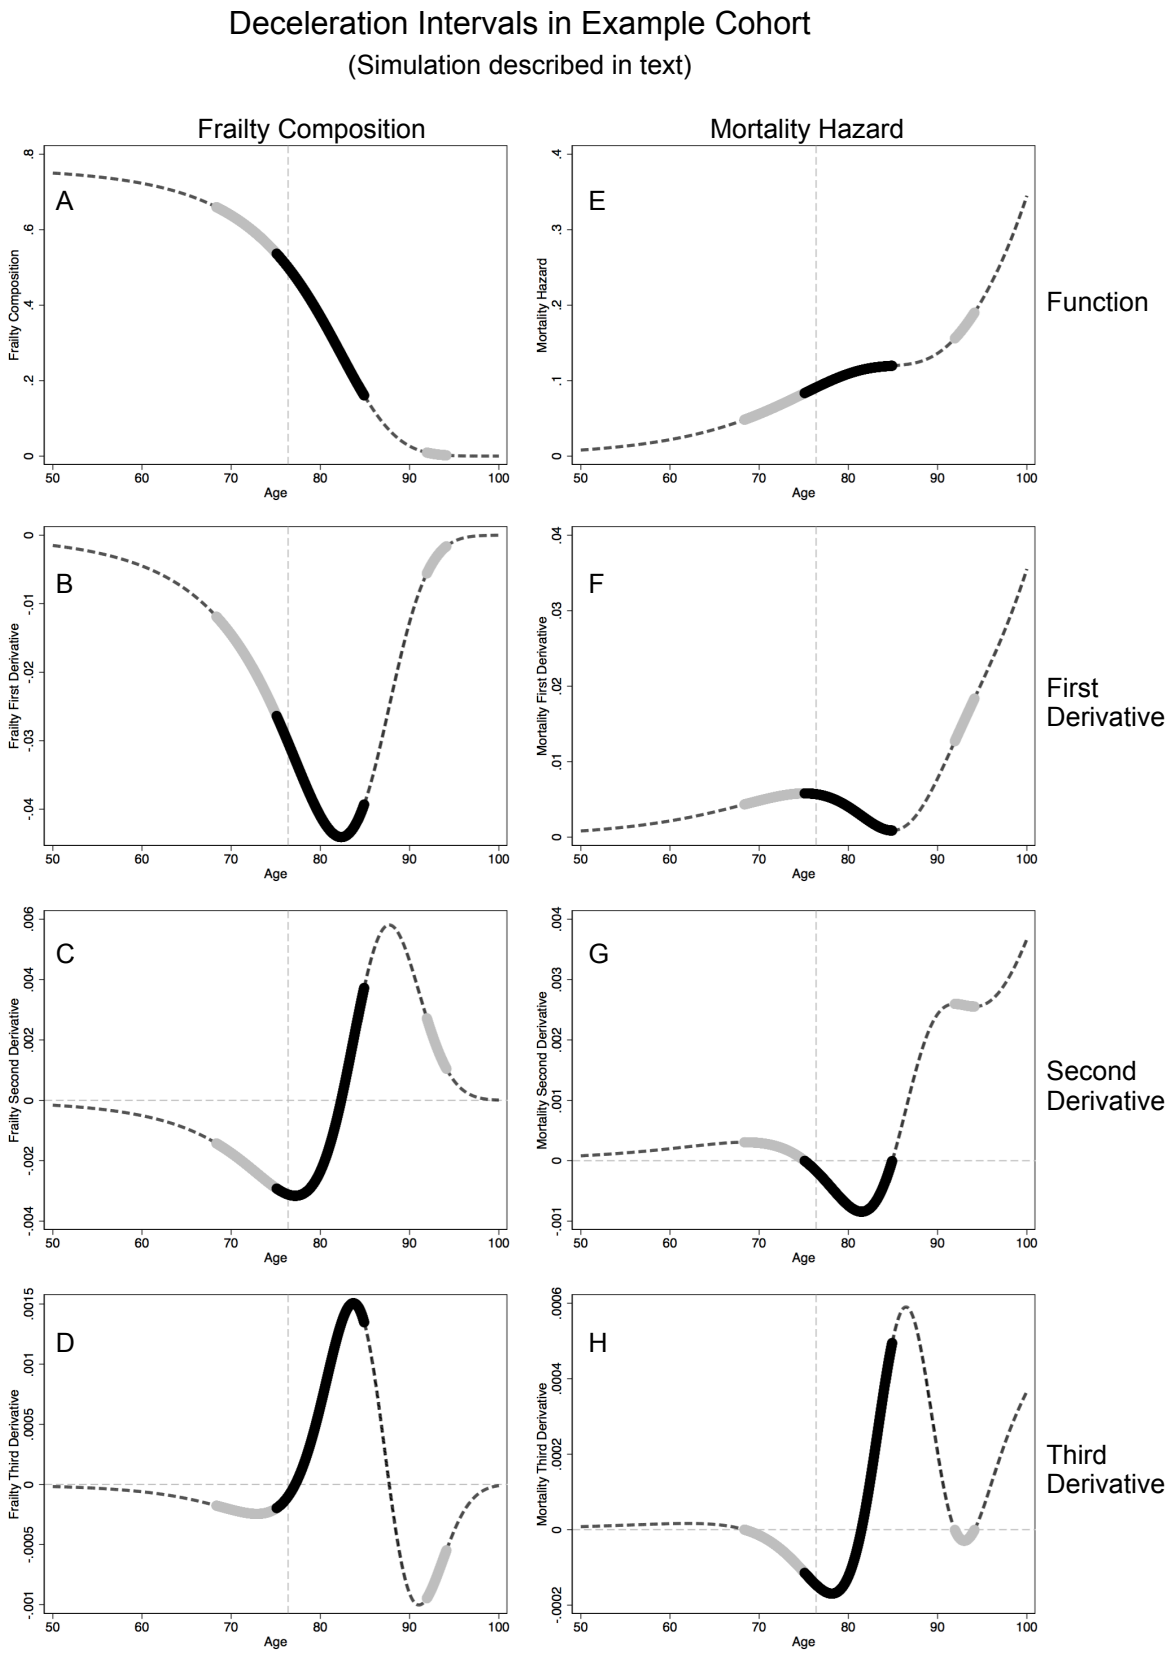
Deceleration Intervals in Example Cohort
(Simulation described in text)
Frailty Composition
Mortality Hazard
A
E
B
F
C
G
D
H
Frailty Composition
Mortality Hazard
Frailty First Derivative
Mortality First Derivative
Frailty Second Derivative
Mortality Second Derivative
Frailty Third Derivative
Mortality Third Derivative
Age
Function
First Derivative
Second Derivative
Third Derivative

Figure 1. Example cohort. The left column gives frailty composition (proportion frail) and the right column gives mortality, both over age. The dashed dark lines represent Gompertz mortality; the thick grey lines, absolute deceleration; and the thick black lines, relative deceleration. The dashed light vertical line marks the point where the frail become a minority, and the dashed light horizontal line in the panels showing the second and third derivatives marks zero.

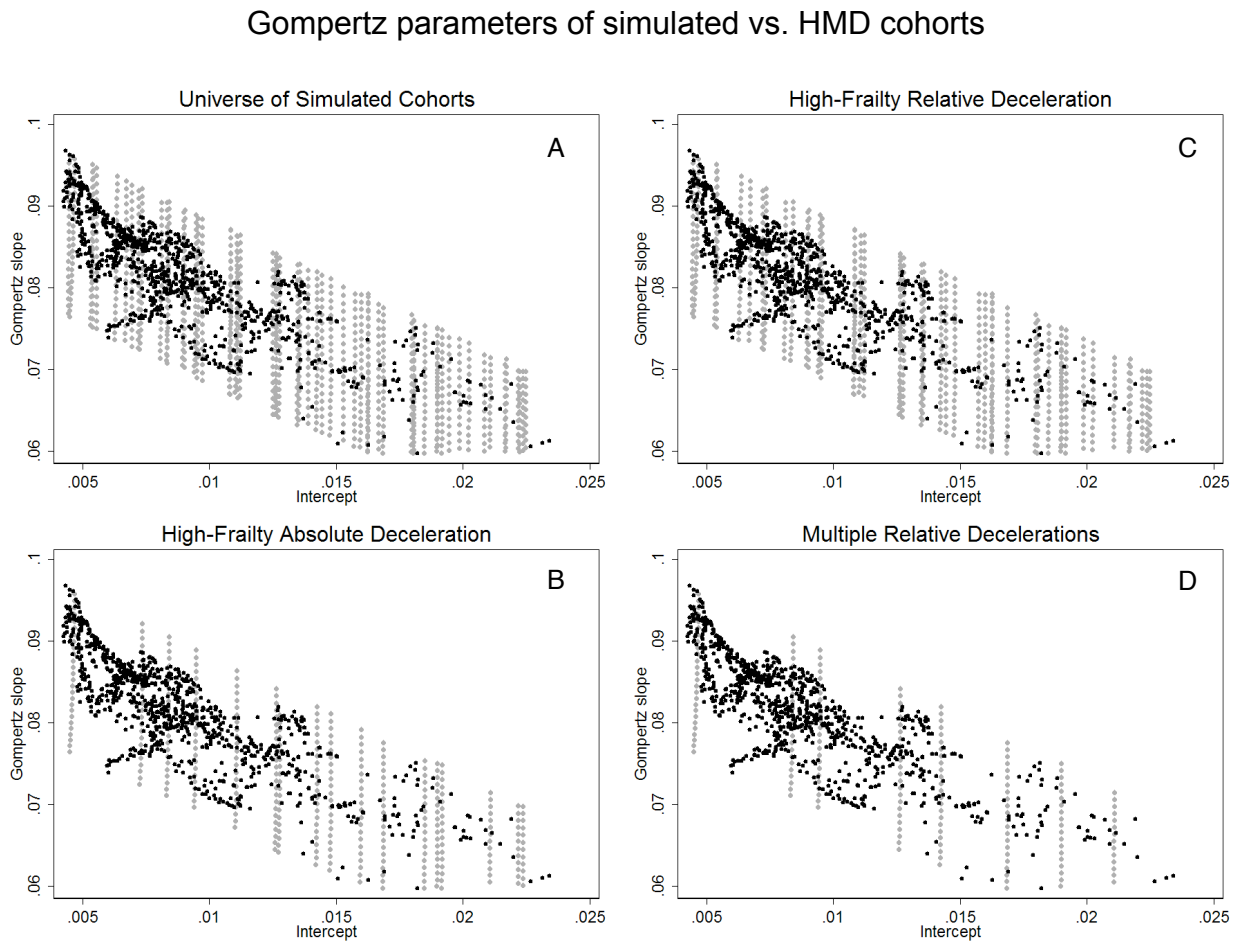

Figure 2. Gompertz parameters of simulated cohorts (grey) compared to Human Mortality Database Cohorts (black). (A) Universe of all simulated cohorts. (B) Simulated cohorts with high-frailty absolute deceleration (absolute deceleration while most of the cohort is frail). (C) Simulated cohorts with high-frailty relative deceleration. (D) Simulated cohorts with multiple relative decelerations.

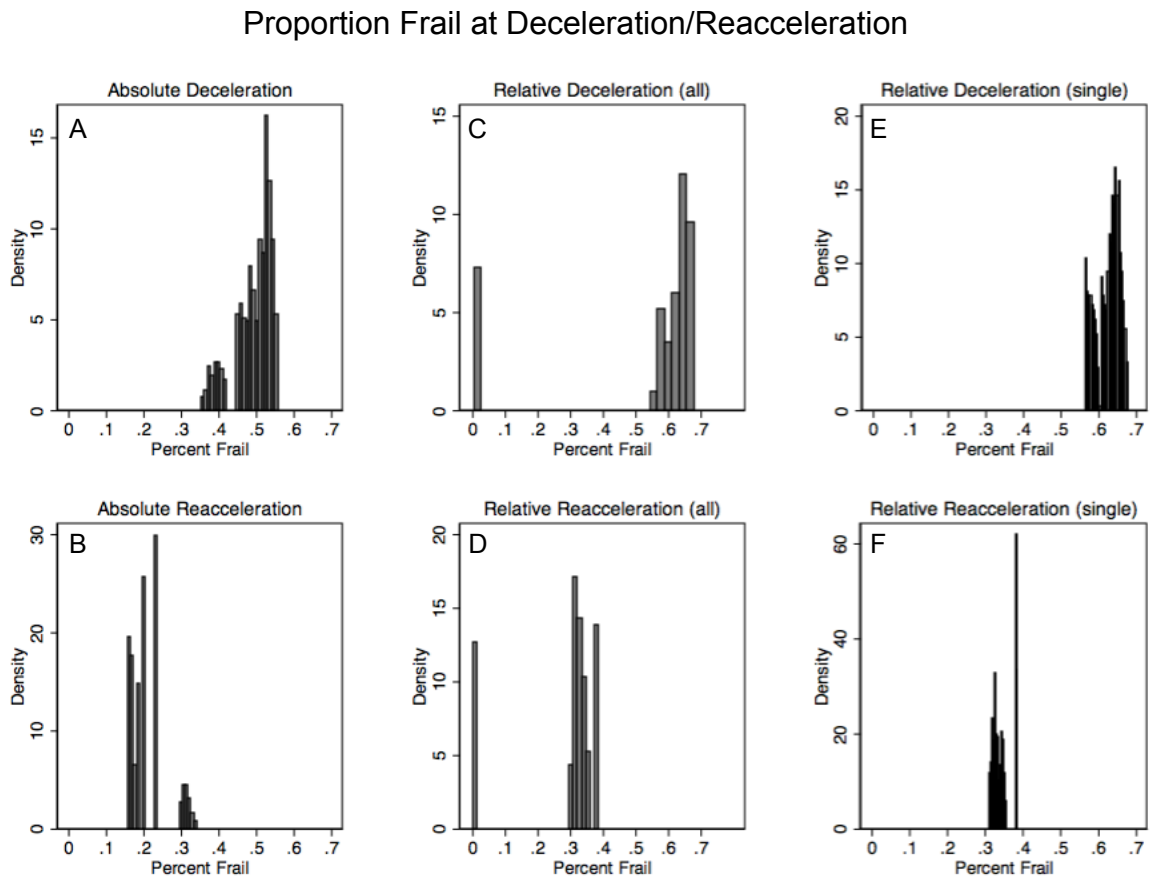

Figure 3. Frailty composition (proportion frail) at deceleration and reacceleration. The top row is deceleration, and the bottom reacceleration; the columns are, respectively, absolute deceleration, all relative decelerations, and relative decelerations limited to cohorts that decelerate only once.

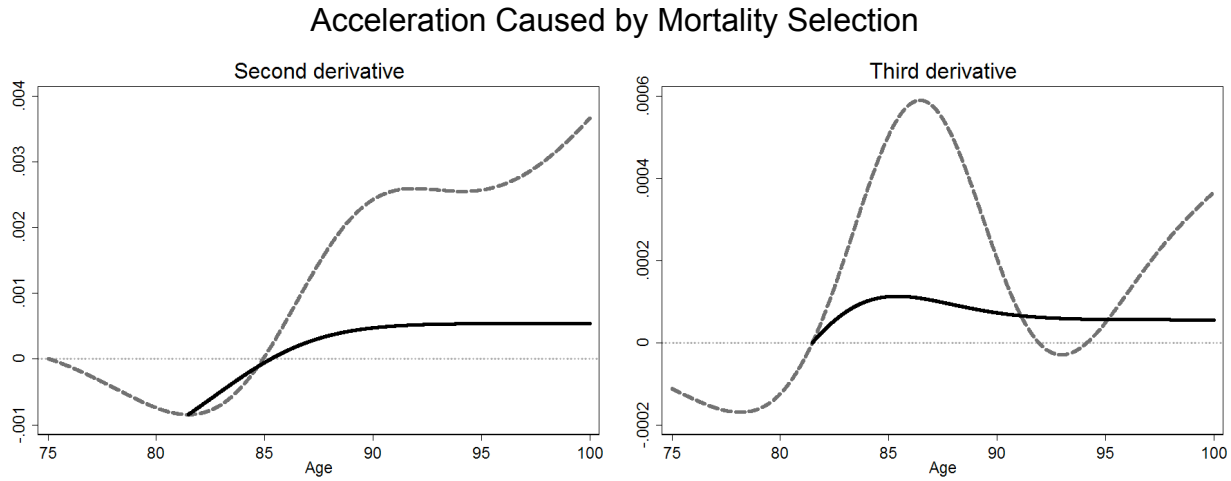

Figure 4. Acceleration caused by declining frailty composition. The solid black line gives the artificial derivatives calculated by fixing subpopulation mortality and allowing frailty composition to

decline as normal. The dashed dark gray line, provided for reference, is the actual derivative of the underlying cohort. The dotted light gray zero line is provided for reference.

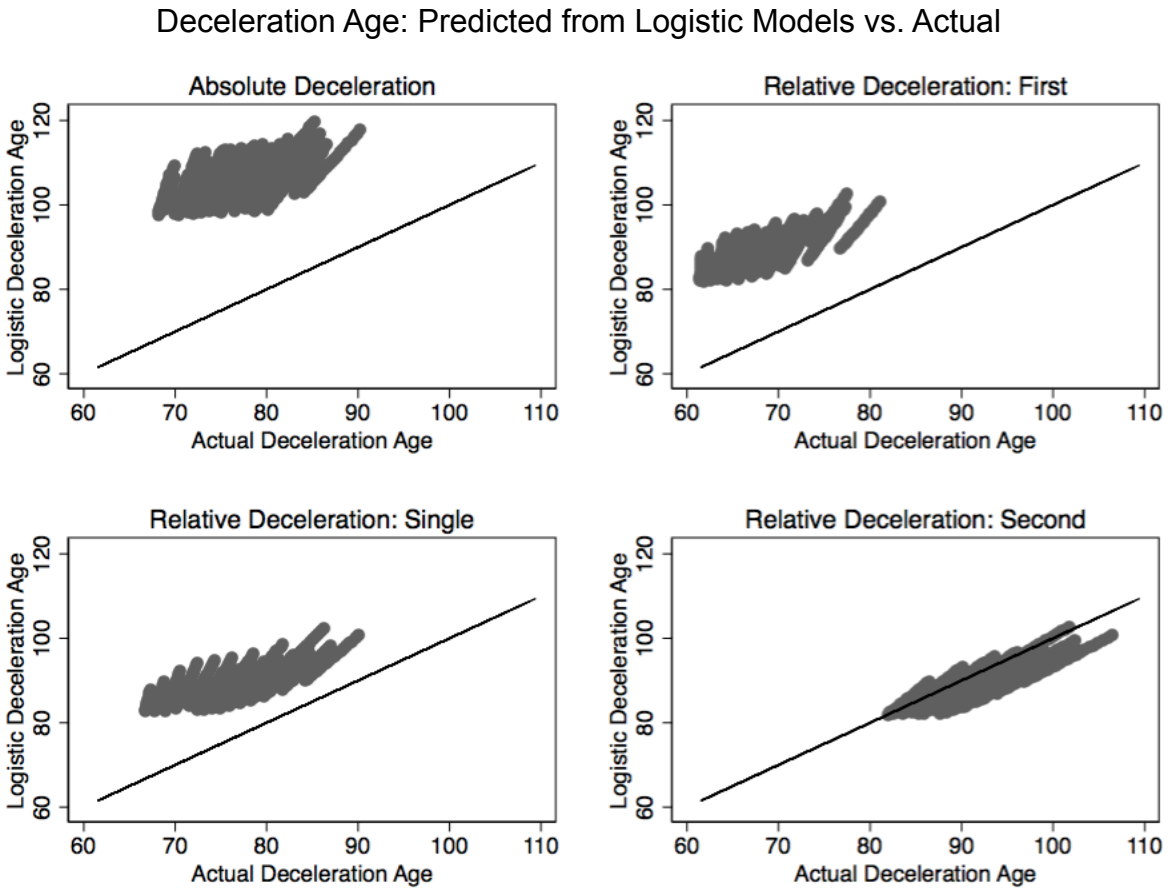

Figure 5. Age at deceleration as calculated from logistic mortality models, compared with the actual age of deceleration in the cohorts. The main diagonal is drawn for reference.

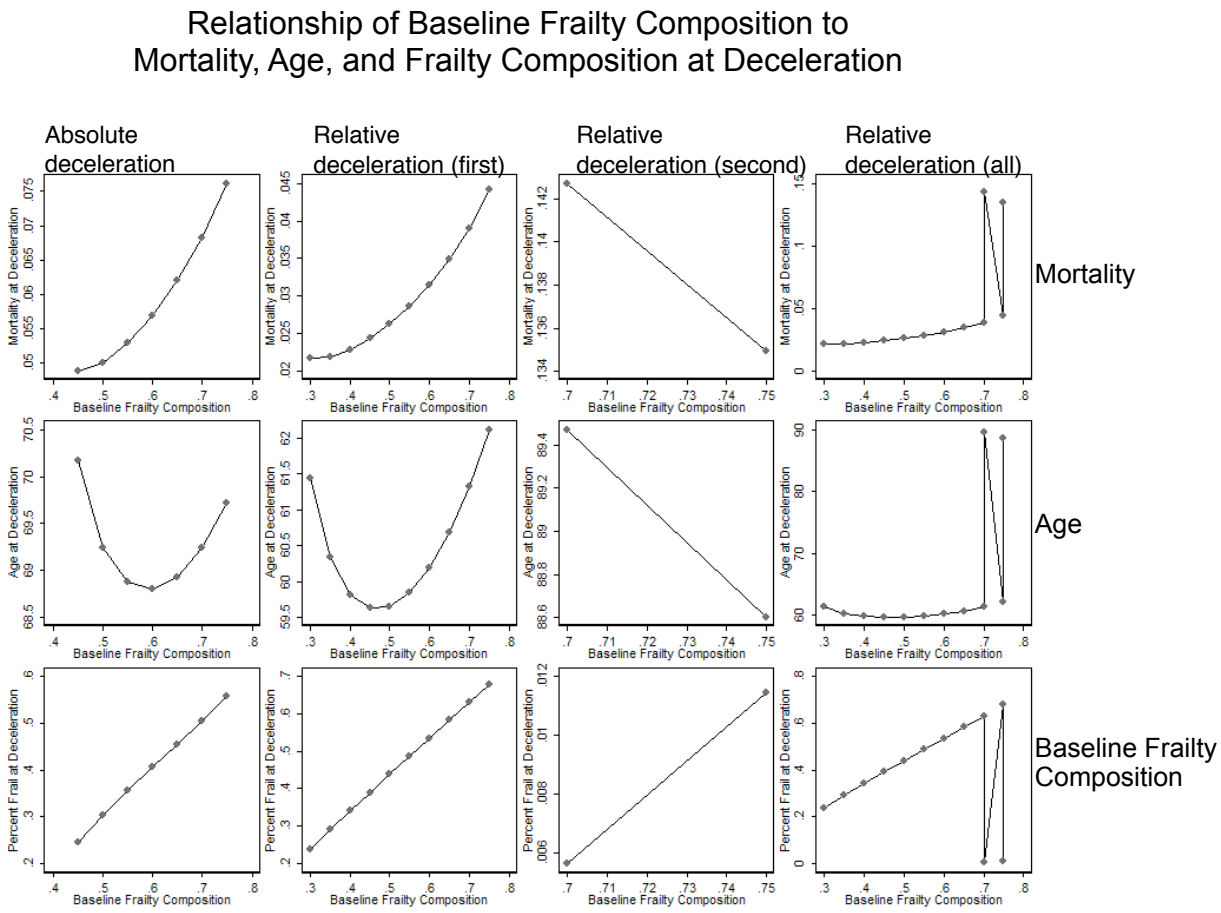

Figure 6. Mortality (top row), age (middle row), and percent frail (bottom row) of deceleration across cohorts with different baseline frailty composition (and fixed other parameters).