What Death Reveals about Life: Mortality Selection and Inequality

By

Elizabeth Wrigley-Field

A dissertation submitted in partial fulfillment of

the requirements for the degree of

Doctor of Philosophy

(Sociology)

at the

UNIVERSITY OF WISCONSIN – MADISON

2014

Date of final oral examination: 7/23/2014

The dissertation is approved by the following members of the Final Oral Committee:

Felix Elwert, Associate Professor, Sociology
Erik Wright, Professor, Sociology
James Montgomery, Professor, Sociology
Jenna Nobles, Assistant Professor, Sociology
Alberto Palloni, Professor, Sociology
James Walker, Professor, Economics

Acknowledgements

Felix
The list of things I must thank you for is multidimensional, dynamic, and far too long for a dissertation acknowledgements page.

Erik
Thank you for showing me what it means to read charitably, to stick to your convictions while you question and listen, to be bold, in short, *to be willing to be wrong*—and therefore to have a chance of being right about something important. (More prosaically, thank you for teaching me to "find the missing fourth" in my implicit 2x2 tables! It led to more than one insight in this dissertation.)

Jenna
Thank you for giving me a vision of what I want demography to be and what I hope I can contribute to it. It made a bigger difference to me than you might realize.

James
Thank you for your generosity with your time and your careful attention, which gave me the help and confidence to stretch my math far enough to do this work. I'm proud of what I learned, and very grateful to you. And I'll stay grateful as I figure out how to stretch further to learn my next things.

Alberto
Thank you for always seeing the big picture, and always being willing to share it with me. It invariably reminds me of how glad I am to be in this tradition with you, demography. I hope I can keep having conversations with you for many years to come—maybe about demography and the microbiome?

Jim
Thank you for encouraging my interest in the crossover very early on, and for welcoming me as a latecomer into CDE. I'm so glad that I became a part of that community.

Writing Support
Thank you to my writing group: Jeff Grigg, Paul Hanselman, Anna Haskins; and thank you to the Writing Center, especially Michelle Niemann and Jenny Conrad.

Funding
Thank you to the Department of Sociology, the Center for Demography and Ecology and Center for Demography of Health and Aging, and the Institute for Research on Poverty at UW-Madison; and to the Ford Foundation and the National Science Foundation for graduate fellowships.

ABSTRACT

A major reason to study disparities in mortality is that they seem to offer evidence about social inequality. Mortality selection offers a generalized threat to such inferences. Within any social group, mortality is not random: some people are consistently more vulnerable to death than others, so as each cohort ages, it contains ever-fewer of these more vulnerable people. Moreover, this selection happens more quickly among disadvantaged, higher-mortality groups. The result is that, at older ages, disparities in mortality—or factors associated with mortality—compare more-intensely and less-intensely selected groups. Unequal survival biases inequality measures.

This dissertation reconsiders two phenomena sometimes thought to shed light on mortality selection: mortality deceleration and the black-white mortality crossover. Mortality deceleration is the slowing of mortality's rise with age, which may result from the progressive loss of 'frailer' members. The crossover is the phenomenon that, at old ages, black mortality falls below white mortality. This crossover may occur because blacks, with higher mortality than whites earlier in life, are thereby more sharply selected. I argue that neither phenomenon offers clean evidence of selection patterns or of underlying social inequalities.

First, I show that a common set of intuitions about mortality deceleration is wrong. Contrary to expectation, even with just two subpopulations with exponentially accelerating mortality, mortality can decelerate twice; it can decelerate while most of the cohort is 'frail'; and selection can produce acceleration in mortality, not just deceleration. These possibilities reflect an unexpected age patterning in the rate of selection. The results imply that standard models may misestimate deceleration timing by decades, and refute some presumed links between deceleration and inequality patterns.

Next I consider the empirical basis for concluding that the crossover reflects mortality selection, or conversely, a genuine reduction over age in black disadvantage. Recent work increasingly imagines selection acting on complex, multidimensional heterogeneity, but derives predictions from simpler unidimensional models. I show that multidimensional selection does not generally support the predictions of simpler models.

My results generally suggest that mortality artifacts provide little evidence for or against specific selection models, and therefore, little evidence about inequality.

TABLE OF CONTENTS

Introduction

Understanding inequality has always been a central demographic motivation for studying mortality (alongside such other purposes as population projection). This means that mortality selection is a generalized threat to many kinds of inferences about social inequality. This is clearest when the evidence of inequality is disparities in mortality itself: we want to use those disparities in mortality to understand something about the unequal conditions of life that produce them, but, except at the youngest ages, populations' mortality disparities conflate those differences in current conditions with the legacy of past inequalities. That legacy is reflected in selective survival: only the strongest survive to old age, and the more disadvantaged in general survival a population is, the truer that is. One result is that, in comparing higher-mortality groups (say, blacks) to lower-mortality groups (say, whites), at ages after lots of mortality has occurred already, we are comparing the hardiest, most mortality-resistant, most *selected* blacks to a much broader cross-section of whites. As a measure of current inequality, this comparison is inherently biased. For example, at the extreme, the group that has experienced far greater disadvantage throughout life may actually have lower mortality late in life than the continually advantaged group, due solely to mortality selection.

Moreover, the problem of mortality selection afflicts not only mortality itself as evidence, but disparities in any measured characteristics that are correlated with mortality as well—which is nearly everything. For example, any study of, say, the effects of one's parents' social class on one's own eventual chance of being able to comfortably retire will face a mortality selection problem: if those raised by working-class parents are less likely to survive to normative retirement age than those raised by richer or higher-status parents, then the ones who

nevertheless do live to such an old age may not be a random sample of the original cohort of working-class kids. These survivors may be the working-class kids who faced advantages in some way that we did not measure, and those advantages may aid their prospects for retirement as well. Here, too, mortality selection would lead us to understate inequality. The far-flung nature of the stylized example makes the point: mortality selection is everywhere. Demographers must devise ways to understand what mortality selection is doing to the disparities we measure, so that we can extract from those disparities the information about inequality that we seek.

This dissertation reconsiders two primary means that demographers have devised to try to extract from mortality patterns information about how mortality selection works: mortality deceleration and the black-white mortality crossover. Each comprises half of the dissertation. Mortality deceleration is the slowing of mortality's rise with age. This deceleration may result from the progressive loss of 'frailer' members as each cohort ages. The mortality crossover is the phenomenon that, at old ages (around age 85 for recent cohorts), black mortality falls below white mortality. This crossover may result from the fact that blacks, with higher mortality than whites earlier in life, are thereby more sharply selected. The hope has been that, by using these presumed artifacts of selection to understand how selection works, we might gain the knowledge to rescue the broad range of inequality analyses that are threatened by mortality selection. Whereas in the ordinary course of disparities research, mortality selection hides inequality, the hope has been to use the mortality phenomena that most clearly suggest selection's operation to reveal inequality instead.

The main argument of this dissertation is that neither of these phenomena—mortality deceleration and the black-white mortality crossover—offers clean evidence of the workings of

selection. This fact renders it an urgent problem of demographic theory to find alternative ways of understanding, documenting, and compensating for mortality selection. This argument has two foundations in the dissertation, two papers presenting the thrust of my views on mortality selection, bookending two smaller papers that flesh out the vertebrae provided in the first and last chapters.

The first paper, "Mortality Deceleration and Mortality Selection: Three Unexpected Implications of a Simple Model," shows that a common set of intuitions about mortality deceleration is wrong. The intuitions go awry because of an unexpected mathematical fact about mortality selection: all else equal, selection happens most quickly when the "frail" and "robust" subpopulations are close to equal size. The purpose of understanding mortality deceleration, if it is indeed driven by mortality selection, is to use its age patterning to understand the patterning of hidden heterogeneities. This enterprise requires that we get two things right: our estimates of deceleration age, and our theory about how differences in deceleration age relate to inequality. I show that a consequence of this unexpected dynamic in selection models is to undermine both steps: standard parametric mortality models can dramatically bias estimates of the age at deceleration, and there is little predictable relationship between, on the one hand, the relative size of the frail subpopulation in one population compared to another, and on the other, those populations' order of deceleration.

The first paper's simple model left standing one generalization about mortality deceleration order: all else equal, higher-mortality groups should decelerate at younger ages. The second paper, "Mortality Deceleration is Not Informative About Unobserved Heterogeneity in Open-Group Settings," shows that this need not be the case when the groups in question can be entered

as well as exited (or can be exited by means other than mortality). In the course of showing this, I develop a model in which the hidden heterogeneity that structures mortality risk ("frailty") also influences the movement between groups (envisioned as health statuses)—a realistic assumption for many groups of interest to social scientists.

The upshot of these two mortality deceleration papers is to suggest that, even under the most favorable assumptions (such as: that mortality deceleration does indeed stem entirely from mortality selection), mortality deceleration patterns offer exceedingly limited evidence about population heterogeneity. In our quest to understand mortality selection so that ultimately we can compensate for its consequences in our estimates of old-age inequalities, mortality deceleration is probably not the artifact we have been seeking.

The other two papers deal with the black-white mortality crossover. The third paper, "Mortality Selection or Life Course Changes? The Missing Explanation for the Black-White Mortality Crossover in the United States," points out that most recent demographic work has assumed, rather than shown, that the crossover (if it is not an artifact of African-American age misreporting in data sources) stems from mortality selection. That there could be a genuine reduction in black survival disadvantage at old ages, relative to younger ages, is no longer even formulated as a hypothetical possibility in much recent work. I show that a perverse consequence of this unstated assumption is that, since everything that might produce a mortality crossover is assumed to be an example of mortality selection, mortality selection is itself misdescribed. This paper is not a call to believe that black disadvantage in survival does decline over age, but it is a call to admit the question and to test mortality selection models rather than choosing in advance when to assume that they are the main source of demographic patterns.

The final paper, "The Selection Problem for Testing the Selection Explanation of the Black-White Mortality Crossover," builds on that call by critically evaluating an increasingly prominent strategy for testing selection models: estimating the age at crossover conditional on some observed dimension of heterogeneity—some type of social inequality—that is hypothesized to play a role in creating the crossover. I argue that this method implicitly draws on a selection model whose workings are far more complex (and substantively realistic) than the model that has been comprehensively analyzed in the theoretical work on mortality selection. But, as I show, the more complex and realistic model does not work in the same way as the simpler models, and does not support the inferences that have been drawn in recent empirical work. The conclusion I draw is that, when the hidden heterogeneity within black and white subpopulations can be assumed to be multidimensional and dynamic—as, indeed, I think it always can be—then the age at crossover, conditional on any particular dimension of heterogeneity, cannot serve as a test of the mortality selection model. The age at mortality crossover, like the age at mortality deceleration, ultimately cannot offer the evidence we are looking for.

* * * * *

There is a perspective that runs through these four papers. These are its three main elements:

First, theories should be testable and tested to the greatest extent possible, and for that, they must be precise—and so must be the alternatives to them.

Second, the deepest contribution of demographic methods is the way they let us shift perspectives between less and more aggregated groupings, between individuals, subpopulations, and populations. That contribution is reflected in demography's having clearly posed the question about mortality selection (which afflicts every social science) in the first place—imagining the heterogeneities hiding in the aggregate death rates, and reasoning about how mortality would change their shape.

In these papers, the power of this demographic approach is reflected in the first, third, and fourth papers, when I show that formulating selection models more explicitly reveals that looser, verbal statements in some prior work have misled even demographers about how mortality selection works. This approach—of moving back and forth between full populations, subpopulations, and individuals to understand how individual-level status transitions are reflected in sometimes-surprising aggregate patterns—is also reflected in the extensions I offer to the classical models, in the second and fourth papers. These papers ask what we should expect to see in subpopulations that are defined by inequalities (e.g., in health or in poverty) that create inequality in mortality, but are not the only inequalities that do so.

A specific instance of that "perspective-shifting" contribution, and arguably the greatest conceptual contribution of demography, is the concept of the cohort: the concept that lets us link a life course perspective or biological theories of aging to the consequences of historical change. Cohorts are the unit that mortality selection acts upon, and are the basis of all of the theoretical models in the paper.

Third, social stratification is complex and multidimensional, and this creates both enduring, "clumpy" inequalities, and cross-cutting, intersectional ones; inequalities that are stable and

entrenched, and others that change their shape as people age. I take up the nature of stratification, as represented in demographic selection models, in each paper.

In the first paper, "Mortality Deceleration and Mortality Selection: Three Unexpected Implications of a Simple Model," discussion of competing views of stratification arises in the discussion of binary vs. continuous, unimodal frailty models, where I compare the kind of model (binary frailty) used throughout this dissertation to another used in some demographic work. The second paper, "Mortality Deceleration is Not Informative About Unobserved Heterogeneity in Open-Group Settings," is motivated by the desire to understand how selection patterns change when the characteristics that define the groups we compare are simultaneously causes of mortality and themselves caused by other inequalities that can kill.

The third paper, "Mortality Selection or Life Course Changes? The Missing Explanation for the Black-White Mortality Crossover in the United States," has a section ("What is mortality selection about?") discussing the substantive meaning of what I identify as the key assumption of all mortality selection models, the "stable ordering" of individual mortality over age. In short, mortality selection is a theory about inequalities that persist over the entirety of the life course. And the central theme of the final paper, "The Selection Problem for Testing the Selection Explanation of the Black-White Mortality Crossover," is extending mortality selection models to encompass more substantively realistic, multidimensional, dynamic models of inequality in mortality risk.

In short, the task is to unite the precision and clarity of formal demography with the substantive insight of the sociology of social stratification. The papers in this dissertation motivate this task (in part by showing that, in the absence of theoretical clarity, substantive

conclusions can be seriously mistaken), but they also suggest that it will not be easy: I show that our tools for understanding mortality selection may be uninformative in the face of more substantively realistic assumptions about population heterogeneity than mortality selection models have traditionally made.

I see this project as connected to a larger view of the status of demography as a field. Demography is a very different field than it was a generation ago. It is more interdisciplinary and concerned with a far broader array of population processes than the traditional three (births, deaths, migration): marriage, divorce, and cohabitation; health, broadly construed; education; labor force participation; and incarceration are now all major topics in the most traditional of demographic venues. Being a demographer no longer requires either working on traditional demographic topics or working with traditional demographic models.

The result is a gap between successive intellectual generations in demography, and a missed opportunity. Where the older demographic approaches excelled at clarity and precision, linking population patterns with immense creativity, those same classics sometimes suffered from a lack of substantive realism. This can be seen in the early assumption that mortality selection was really a story about genetics, not about social inequality; or in the debate that formed the backdrop to much of the early mortality selection research (especially on mortality deceleration), on whether there is an inherent limit to human lifespans, that said nary a word about how social arrangements might constrain the answer to that question. The new demography has shed this narrowness. It cares about inequality of all kinds, and it assumes that they should be understood together. But in losing its connection to the older approaches, it is also giving up a great deal of insight. The mortality selection work before roughly 2000 is careful,

creative, correct, and constrained; the empirical work after that is historical, intersectional, and, often, mathematically mistaken. The hope of my dissertation is that the next generation of demography should meld the strengths of each of these.

CHAPTER 1

Mortality Deceleration and Mortality Selection:

Three unexpected implications of a simple model

CHAPTER ABSTRACT

Unobserved heterogeneity in mortality risk is pervasive and consequential. Mortality deceleration—the slowing of mortality's rise with age—has been considered an important window into heterogeneity that otherwise might be impossible to explore. This paper argues that deceleration patterns may reveal surprisingly little about the heterogeneity that putatively produces them. 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 cohort 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 cohorts that in the aggregate resemble cohorts in the Human Mortality Database. I argue that these results: challenge some conventional heuristics for understanding the relationship between selection and deceleration; undermine certain inferences from deceleration timing to patterns of social inequality; and imply that standard parametric models, assumed to plateau at most once, may sometimes badly misestimate deceleration timing—even by decades.

Introduction

Unobserved heterogeneity in mortality is pervasive and consequential. All cohorts are heterogeneous in mortality risk in ways that are often stably ordered over the life course: if one person is higher-risk than another at young ages, they will also be higher-risk at older ages, should they both live that long. This heterogeneity gives rise to mortality selection: the frailest members of a cohort disproportionately succumb to mortality, eventually leaving an intensely selected, relatively robust cohort in its place (e.g., Beard 1959, 1971; Kannisto 1992; Vaupel, Manton, and Stallard 1979; Vaupel and Yashin 1985). This selection will in general occur unequally across social groups with different patterns of heterogeneity and overall levels of mortality risk. Such unequal selection substantially complicates efforts to understand inequalities in mortality risk—including inequalities driven by observable factors—across and within groups.

The challenge posed by mortality selection for accurately modeling differences in mortality arises in part because—unlike in a linear regression context—mortality selection can bias hazard models even when the omitted sources of mortality risk are unrelated at baseline to the covariates included in the model (Rodriguez 1994). There are many examples of the distorting effects of unobserved heterogeneity caused by mortality selection. Perhaps the most well-known is the black-white mortality crossover (e.g., Berkman et al. 1989, Dupre et al. 1996, Kestenbaum 1992, Lynch et al. 2003, Manton et al. 1979, Masters 2012)—one of many documented mortality crossovers (e.g., Rogers 2002)—which confounds racial inequalities operating at old ages with the intensified selection for robustness experienced by African-American cohorts. Other examples abound. A long-running demographic controversy is whether the well-documented narrowing over age of the mortality disparities by educational status reflects a true convergence

of mortality trajectories—perhaps implying a biological limit to lifespan that even healthy behaviors and other advantages cannot overcome—or simply the operation of mortality selection among the higher-mortality less educated, masking a protective effect of education that continues into old age (e.g., Robert and House 2000, Zajacova et al. 2009). Assessing demographic theories about the effects of early-life health on later-life outcomes can also be confounded by mortality selection: the cost to individuals of negative childhood conditions can be masked by the robust composition of those who survive such conditions (e.g., Costa 2012, Palloni 2006). Unobserved heterogeneity can be deeply distorting to efforts to understand group-level inequalities and individual determinants of mortality.

The significance of unobserved heterogeneity in disguising how mortality inequalities unfold over a life course stands in stark contrast to the relative dearth of methods for exploring such heterogeneity. Without comprehensive individual-level measures of the determinants of mortality, population scientists can consider only the aggregate mortality of groups, aided by either parametric assumptions about the unobserved distribution of risk within those groups (in a long research tradition beginning with Vaupel et al. [1979]) or qualitative assumptions about the distribution across groups. Simply put, when dealing with a phenomenon that is by definition unobserved, options are limited.

In this context, mortality deceleration has been considered an important window into heterogeneity that otherwise might be impossible to explore. Mortality deceleration—the slowing of mortality's rise with age—can occur as cohorts are systematically selected by mortality to disproportionately contain those most robust to death (e.g., Beard 1959, 1971; Kannisto 1992; Thatcher et al. 1998; Vaupel et al. 1979; Vaupel and Yashin 1985). Thus,

deceleration has often been considered an important indicator of substantial prior selection—and thus of heterogeneity in the original cohort. Similarly, differential deceleration patterns across groups (e.g., cohorts or race and sex groups) has been used as evidence of differences between groups' patterns of internal heterogeneity (e.g., Horiuchi and Wilmoth 1997, 1998; Lynch and Brown 2001; Lynch et al. 2003). Indeed, 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."

Demographers have, however, also recognized that deceleration patterns are no panacea for understanding heterogeneity since deceleration in mortality risk may arise biologically at the individual level, not only from mortality selection at the group level. In particular, many experimental biodemographic studies have not found heterogeneity alone to be a plausible cause of observed deceleration in insect and animal populations (e.g., Carey et al. 1992, Curtsinger et al. 1992, Drapeau et al. 2000 [but see Steinsaltz 2005], Rauser et al. 2005, Vaupel and Carey 1993), and a smaller number of studies have reached similar conclusions for human data (Mueller et al. 2011, Steinsaltz and Wachter 2006; and see reviews in Vaupel 1997, Wachter and Finch 1997). If deceleration does arise for reasons other than selection, then heterogeneity might or might not interact with individual-level decelerating mortality to influence population-level rates, but deceleration would not constitute straightforward evidence of selection (see Steinsaltz and Evans 2004 for an elaboration of the argument that deceleration patterns alone are not telling evidence for any particular model that might give rise to them).

In this paper, I show that *even if* deceleration stems entirely from mortality selection on unobserved heterogeneity, those deceleration patterns still may reveal little about the

heterogeneity that produced them. The reason is that even an exceedingly simple binary frailty model can produce deceleration and acceleration patterns sufficiently complex so as to defy some conventional models and predictions. In particular, I show that mortality can decelerate even while most of the cohort remains frail; that mortality can decelerate twice, even with only two subpopulations; and that mortality can reaccelerate, not only because mortality selection has already run its course (a possibility discussed by Vaupel and Yashin [1985]), but also—counter-intuitively—because of the continued action of selection itself. These possibilities stem from a simple mathematical result: the rate of selection is greatest, *ceteris paribus*, when the frail and robust are each half of the cohort. Thus, this paper explores the relationship between the level of frailty in a cohort as it ages, its rate of selection, and the deceleration and acceleration patterns produced by those two things.

These results have two additional implications. As I show, parametric mortality models, such as logistic models, may badly misestimate deceleration timing—even by decades. And standard cross-cohort comparisons of deceleration timing may misstate the direction of inequality in mortality risk.

I begin by stating more explicitly some conventional expectations about mortality deceleration that these results call into question. For concreteness, I then provide an example of the kind of counter-intuitive pattern that the analysis shows to be possible, before presenting the key analytical result that provides the intuition for the mortality patterns explored. However, the complexity of the mortality derivatives makes it difficult to assess analytically when such counter-intuitive patterns will arise, so the main body of the paper presents the results of simulation models with parameters drawn from the Human Mortality Database. Simulations

show that the surprising patterns described here can be pervasive in models with realistic parameters, and that the resulting problems with parametric mortality models and certain cross-group comparisons may also be widespread. Before concluding, I argue that, since the results explored here are precluded by widely-used gamma-distributed frailty models, these results suggest a need for more direct comparison of heterogeneity models.

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] compares 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

[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).

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, or may stop or reverse at any point. This paper is agnostic about which of absolute or relative deceleration, if either, is a preferable measure, and explores the properties of both. Accordingly, the crucial measures defining deceleration and reacceleration in what follows will be the signs of the derivatives of mortality.[2]

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

[2] 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; Horiuchi and Coale 1990; Horiuchi and Wilmoth 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 are functions only of the derivatives of cohort frailty composition and of subpopulation mortality (as shown in Appendix 1). Using the derivatives of mortality in its own scale therefore allows us to more cleanly focus on the contribution of mortality selection, that is, of the declining composition of frail members of the cohort.

> 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][3]

Heathcote, Puza, and Roberts (2009) makes a similar claim in a paper giving rare explicit 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 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—at least if there are only two 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 dece lerations are precluded in a two-subpopulation model is made explicit in

[3] 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]

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.[4]

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

[4] 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 subpopulations can experience two successive decelerations.

[5] Demographic intuitions on this point may be influenced by a result presented in perhaps the most influential paper introducing mortality selection to a wide demographic audience, Vaupel and Yashin's (1985) "Heterogeneity's Ruses." Vaupel and Yashin (1985: 177) write:

> The sudden decline in the observed hazard rate is produced by the rapid extinction of the frailer subcohort. Until the point of decline, the frailer subcohort experiences death rates that are relatively low. Then, due to the exponential increase in the force of mortality, the death rates become sufficiently large that within a few years almost all of the frailer subcohort dies. The observed hazard rate declines to the level of the hazard rate for the more robust subcohort. Since this hazard rate is increasing, the observed hazard rate then starts to increase as well: the observed hazard rate now equals the hazard rate for the more robust subcohort because only members of the more robust subcohort are still alive.

Vaupel and Yashin are describing a cohort whose mortality increases, decreases, then increases, rather than the acceleration and deceleration of such a cohort. But the vivid imagery of decline precipitated by the rapid extinction of the frail, and then rising with the mortality of the robust, may have been naturally extended by analysts from hazard slopes to higher derivatives. As we will see, the analogue of Vaupel and Yashin's slope pattern in the third derivative of the hazard is one form of deceleration and acceleration that can occur, but only one form—and not at all the most common form in the model cohorts to be considered here.

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 subcohorts, each with Gompertz mortality, with 75% of the cohort in the frail subpopulation at age 50.[6]

Fig. 1 displays this example cohort from ages 50-100 (by which age the frail are virtually extinct). Panel A of Fig. 1, 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 increasingly accelerates until age 68, when the first interval of relative deceleration begins, with 66% of the cohort frail. At age 75, with 54% of the cohort frail, mortality decelerates absolutely; the second derivative remains negative until age 84 (16% frail), when a second period of increasingly accelerating mortality begins. This persists until age 91 (0.9% frail), when mortality again decelerates relatively until age 94 (0.2% frail), when the cohort enters increasingly accelerating mortality for the third and final time. Panel B shows the slope of the percent frail of the cohort. We will see that this slope—which, with its sign reversed, is the rate of frailty decline—drives much of the results to follow. In this cohort, frailty declines fastest at age 82, when the frail are 27% of the cohort. The dashed vertical line in both panels marks the point when the frail are half of the cohort, which will be shown to be a turning point in some of the selection dynamics explored in this paper. (The full set of derivatives of mortality and frailty for this cohort are shown in Appendix 1.)

[6] I give the full parameters of this example, and justify their reasonableness, in Footnote 15, after I describe the simulations.

Analytical Intuition: High-Frailty Deceleration, Multiple Deceleration, and Selection-Driven Acceleration are Possible in Principle

The surprising mortality patterns shown here to be possible arise from the role played in the mortality derivatives by the rate of change in the cohort's% frail. Specifically, it turns out that the level of frailty—the percent frail—plays two distinct roles in producing mortality acceleration and deceleration: a direct role (as already shown in Vaupel and Yashin [1985]), and an indirect role via the rate of selection. Here I show the key equations that provide the intuition; Appendix 1 provides the full analytical results.

The model in this paper assumes that both subcohorts have Gompertz mortality, as shown in Eq. 1:

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

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

The two subcohorts share a log-slope over mortality, β. That log-slope and the baseline level of robust mortality, α, are assumed to be greater than zero. The *frailty multiplier*, f, which is the ratio of frail to robust mortality at any age, is assumed to be greater than 1.

Cohort mortality at any age, and its derivatives, are functions of three parameters: the mortality of the robust, $\mu_r(x)$; the mortality of the frail, $\mu_f(x)$; and the percent of the cohort that is frail (the frailty composition), $\pi(x)$, as shown in Eq. 2:

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

The crucial role in what follows is played by the slope of frailty composition, that is, by the rate at which the cohort is becoming less frail due to mortality selection. Equation 3, which provides the intuition for all results in this paper, shows that this slope is a function of the level of frailty composition and of the difference between frail and robust mortality:

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

All terms in Eq. 3 are non-negative, and so the negative sign means that the slope is always negative: at all ages, mortality selection makes the cohort less frail. The absolute value of this expression can be considered the rate of selection. The binomial variance of the percent frail, $\pi(x)\big(1-\pi(x)\big)$, is greatest when $\pi(x)=.5$. In other words, *all else equal, the rate of frailty decline is greatest when half of the cohort is frail and half robust*. Yet all else is *not* equal: the absolute difference between frail and robust mortality, $\mu_f(x)-\mu_r(x)$, is greatest when the mortality of each subpopulation is greatest: at the oldest ages.

These two terms create the possibility that selection intensifies, and mortality decelerates, as $\pi(x)$ hurtles downward toward .5; that selection attenuates, and mortality reaccelerates, as frailty declines further away from .5; and that selection reintensifies, and mortality decelerates

again, at very old ages, when $\mu_f(x) - \mu_r(x)$ is very large (before mortality accelerates a final time with only the robust left alive). Yet the mortality derivatives are sufficiently complex that it is difficult to evaluate analytically whether to expect these patterns in cohorts with realistic parameters. To answer that question, 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

Simulation Procedure

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. To minimize the role played by assumptions about unobserved parameters and maximize the role played by real data in the selection of simulations, I generate many candidate latent subpopulation models and keep only those whose aggregate parameters are consistent with the life tables of known cohorts.

The mortality hazards are calculated analytically, in instantaneous time. To limit the complexity over four dimensions, one parameter—baseline frailty composition—is set in all simulations at .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.[7] The model therefore assumes that 75% of the population surviving to age 50 is frail.[8] Thus, these simulations represent cohorts in which mortality advantage, rather than disadvantage, is the exceptional condition.[9] The frailty 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

[7] 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 logistic) models, used for specific purposes described below, 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).

[8] One might be concerned that it is impossible for a cohort to be 75% 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.

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

Some demographic theory on mortality compression suggests that such populations—where longevity relative to one's cohort is the exception rather than the rule—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. 2012). 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 in this paper 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 et al. 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.

extreme of what we might consider plausible for human populations.[10] 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.[11]

To winnow these 640,000 simulated cohorts down to a realistic subset, I estimate a Gompertz model on each aggregated cohort,[12] and keep only the ones that fall inside a parallelogram formed around the intercept and log-slope parameters estimated from the 2,352 historical European cohorts collected in the Human Mortality Database (HMD).[13] The parallelogram hugs the shape of the HMD cohorts (shown visually in Appendix 2), so that cohorts included in the final simulation universe generally not only have a Gompertz intercept and slope similar to real (but distinct) cohorts; they have a *combination* of intercept and slope

[10] For a rough-and-ready sense of what 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).

[11] MATLAB code is available from the author on request.

[12] 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.

[13] 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). I use all cohort (vs. period) data included in the HMD.

that is similar to at least one real cohort.[14] The result is a universe of 1,151 simulated cohorts analyzed in this paper.[15]

Simulation Results

High-Frailty and Multiple Deceleration—At what percent frail does deceleration and reacceleration occur in this universe? Figure 2 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 demonstrate the problems for the heuristics that mortality decelerates only once (one-third of the cohorts decelerate relatively twice) and only when the frail are nearly depleted. Panel A of Fig. 2 shows that absolute deceleration *never* corresponds, in these simulations, to the latter heuristic: absolute deceleration can occur when the frail are a majority or a minority, but never occurs here when they are less than 35% of the population. As shown in Panel B, reaccelerations following absolute deceleration, likewise, occur well before frailty depletion, when the frail are 15-34% 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

[14] The constraint that simulated cohorts resemble a real *combination* of intercept and slope is relevant because of the well-known negative correlation across cohorts between those two parameters, first noted by Strehler and Mildvan (1960) and still of great demographic interest (e.g., Finkelstein 2012; Zheng, Yang, and Land 2011).

The method used here produces a set of simulated cohorts whose parameters are similar to those of real cohorts because the bulk of the real data fall into a parallelogram-like shape. This is particularly true for simulated cohorts close to relatively recent real cohorts, where the data are less sparse. All patterns discussed in this paper occur across the range of Gompertz intercept and slope values in the HMD universe (also shown visually in Appendix 2), including among cohorts that fall amid dense clusters of real data.

[15] Returning now to the example cohort presented in Fig. 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 cohorts 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.

decelerating when the frail are nearly gone (in these cohorts, when frailty composition ranges between 0.5-1.2%) and reaccelerating shortly thereafter. But Panels E and F demonstrate that this pattern occurs only in the second of two decelerations. These panels are limited to the 67% of simulated cohorts with only one relative deceleration. 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. (An analysis of where in the parameter space high-frailty and multiple decelerations occur is presented in Appendix 2.)

Selection-Driven Acceleration—It seems intuitive to conceive of acceleration/deceleration as a tradeoff between accelerating subpopulations (producing acceleration) and declining frailty composition (producing deceleration). Yet 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 at age 81—the age when the aggregate second derivative reaches its minimum (with 31% of the cohort frail)—while the percent frail is left varying as in the actual cohort. This calculation isolates the effects of the declining percent frail from those of increasing subpopulation hazards, slopes and accelerations. Figure 3 shows the results: while mortality does not reaccelerate nearly as dramatically without the changes in the subpopulation mortality, nevertheless it does reaccelerate. In this exercise, the declining rate of selection, as frailty falls further below half, leads mortality to reaccelerate even

in the absence of subpopulation-level changes. This demonstrates that mortality selection can contribute to acceleration as well as deceleration.

Further Implications

Bias in Estimated Age at Deceleration Using Common Parametric Models

Results so far suggest that, even in an exceedingly simple model, it is possible, indeed plausible, for cohorts 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.

Fig. 4 plots the deceleration timing derived from two-parameter logistic models estimated on the simulated cohorts against the actual deceleration timing of those cohorts, with the main diagonal provided as a reference line.[16] If the logistic models work well, the estimated deceleration ages should fall along (or near) this main diagonal. For most decelerations, they fall very far from it.

Panel A of Fig. 4 shows the results for absolute deceleration. The logistic models badly overestimate the age of deceleration, with the overestimation ranging between 18 and 39 years

[16] Deceleration timing for logistic models is defined in the same way as for the nonparametric true hazards—that is, when the second or third derivative become negative—using formulas for those derivatives taken from Rau et al. (2009).

(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. These panels show that the logistic models fit poorly the deceleration patterns for first and single relative 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 (relatively rare) second relative decelerations by between 1 and 6 years (underestimating by only 3 on average).

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. The 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 results are especially problematic for two reasons. First, the falsely detected decelerations are found at very similar ages to those identified in previous empirical research on the age at deceleration. Second, the magnitude of the error found in the estimation of age at

deceleration—up to several decades—dwarfs the size of differences in deceleration timing interpreted substantively in the empirical literature, which are sometimes only a few years.[17]

It is routine in the demographic literature for logistic or similar parametric models to be estimated on mortality data without checking the raw data for evidence of multiple deceleration, or deceleration followed by reacceleration. This might be because the raw data are too noisy to support such investigation non-parametrically, but perhaps it is also because such deceleration patterns are not considered a serious possibility, or a serious source of model bias if they do exist. The latter two ideas are cast into doubt by these results. These results collectively suggest that it would be advisable for demographers whose primary object of study is deceleration not to rely solely on models that assume a single peak in mortality. It may be best to use these conventional parametric approaches alongside some effort to consider whether multiple deceleration, or deceleration with reacceleration, may be present in the data.

Comparing Deceleration across Cohorts

A central motivation for accurately measuring deceleration is that comparing deceleration timing across cohorts or other closed social groups may permit demographers to infer something

[17] The empirical literature finds deceleration at very similar ages to those identified—sometimes erroneously—here, and reports changes in deceleration timing that are often much smaller than the potential two-decade error identified here. For example, 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; 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 age misreporting), from 1970-1992.

about the differences in the cohorts' distributions of mortality risk (Horiuchi and Wilmoth 1997, 1998; Lynch and Brown 2001; Lynch et al. 2003). This endeavor, which links the measurement of deceleration to the study of inequality and change in mortality, necessitates a qualitatively simple relationship between unobserved patterns of heterogeneity and observed patterns of deceleration. 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.[18] 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). In this section, I show more explicitly how this prediction will sometimes fail.

To test this prediction using all the parameter combinations in the simulation universe, I compare deceleration timing in pairs of cohorts with fixed frailty multiplier f, robust intercept α, and log-slope β, but with baseline frailty composition π_0 varying in units of .05 from .05 to .75. Table 2 reports the proportion of pairs of cohorts for which the mortality or age at deceleration is greater 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. 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

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

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 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.[19] (These results are extended to consider the percent frail at deceleration, and an example set of cohorts is offered to illustrate the results in this section, in Appendix 3.)

Parameterizing Heterogeneity

The results in this paper do not readily generalize to the widely used models with gamma-distributed individual frailty and Gompertz individual hazards, since those models aggregate to a logistic (hence single-deceleration) hazard (Beard 1959, 1971). Such gamma-Gompertz models fit many data well (Steinsaltz and Wachter 2006, Missov and Finkelstein 2011) and have been very widely used (e.g., Gampe 2010, Horiuchi and Wilmoth 1998) ever since Vaupel, Manton, and Stallard (1979) pointed out their convenient properties, but whether they are appropriate for all populations or whether competing models might fit equally well is a question that should receive further attention (Steinsaltz and Evans 2004).

[19] Moreover, in practice, many comparative analyses will compare the mortality of cohorts that differ in their frailty distribution at whatever age is taken as baseline. This is because 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 considerably in their degree of heterogeneity (Edwards and Tuljapurkar 2005). The results in this paper 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.

In particular, it is not clear which model might be preferable in cases of extreme stratification along unobserved dimensions. Gamma-Gompertz models differ from the dichotomous model used here in two respects: the heterogeneity they estimate is continuously distributed, and it is unimodal. In cases where one unobserved dimension of heterogeneity is so extreme that it swamps other sources of variation between individuals, it is not obvious *a priori* whether a continuous but single-peak model is a better approximation than a dichotomous one.

To the best of my knowledge, all previous models either capture continuous individual variation at the cost of imposing a unimodal distribution conditional on observed covariates (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). The results presented here suggest that it may be useful to more directly compare discrete and gamma-Gompertz models, particularly in populations that may have unmeasured extreme, categorical stratification, since the models produce such divergent deceleration patterns. Further innovations in modeling heterogeneity, such as mixtures of gamma distributions—which could combine the potential virtues of continuous variation and multiple peaks—should also be explored in future research, including research into the deceleration and acceleration patterns such models may produce.

Since gamma-Gompertz models preclude the deceleration and reacceleration patterns described here, in principle, these theoretical results suggest a test of those models, by looking for multiple decelerations or reaccelerations in empirical data (analogously to Horiuchi and Coale's [1990] classic work using higher moments to distinguish models generating similar

estimated means). In practice, such a test would require very high quality data at old ages, since deceleration would need to be assessed nonparametrically.

Conclusion

Unobserved heterogeneity is a barrier to accurately modeling individual-level risk and group-level inequalities in many domains of mortality research, and options for investigating it are relatively limited. Mortality deceleration has been considered by many a promising avenue for assessing such heterogeneity. This paper has shown that drawing conclusions about heterogeneity from mortality deceleration is problematic for reasons that have not previously been articulated, by demonstrating 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 mortality selection 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 misestimate—by decades—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 three 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 sometimes decelerate at *lower mortality* and *younger ages* than one with fewer frail members at baseline.

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. These results call for caution in modeling and interpreting mortality deceleration. 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. 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.

Chapter Acknowledgements

This research was supported by a graduate research fellowship from the National Science Foundation, a dissertation fellowship from the Ford Foundation, and core grants to the Center for Demography and Ecology (R24 HD047873) and Center for Demography of Health and Aging (P30 AG017266) at the University of Wisconsin-Madison. The author thanks, for helpful feedback, Felix Elwert, James Montgomery, Jenny Conrad, Michal Engelman, Duncan Gillespie, Jeffrey Grigg, Paul Hanselman, Anna Haskins, Michelle Niemann, Jenna Nobles, and Sarah Zureick-Brown, along with excellent reviewers and the Editor, and, for technical assistance, Russell Dimond, David Siegel, and the Social Science Computing Cooperative at the University of Wisconsin-Madison.

References

Beard, R. E. (1959). Notes on some mathematical mortality models. In Wolstenholme, G.E.W. & O'Connor, M. (Ed.) *The Lifespan of Animals* (pp. 302-311). 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* (pp. 57-68). 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.

Berkman, L., Singer, B., & Manton, K. (1989). Black/white differences in health status and mortality among the elderly. *Demography*, 26(4): 661-678.

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. (2012). "The Significance of Education for Mortality Compression in the United States." *Demography*, 49(3): 819-40.

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.

Costa, D. L. (2012). Scarring and mortality selection among Civil War POWs: A long-term mortality, morbidity, and socioeconomic follow-up. *Demography*, 49(4): 1185-1206.

Curtsinger, J. W., Fukui, H. H., Townsend, D. R. & Vaupel, J. W. (1992). Demography of genotypes: Failure of the limited life-span paradigm in Drosophilia-Melanogaster. *Science,* 258**,** 461-463.

Drapeau, M. D., Gass, E. K., Simison, M. D., Mueller, L. D., Rose, M. R., (2000). Testing the heterogeneity theory of late-life mortality plateaus by using cohorts of Drosophila melanogaster. *Experimental Gerontology*, 35 71-84.

Dupre, M. E., Franzese, A. T. & Parrado, E. A. (2006). Religious attendance and mortality: Implications for the black-white mortality crossover. *Demography,* 43**,** 141-164.

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.

Finkelstein, M. (2012). Discussing the Strehler-Mildvan model of mortality. *Demographic Research*, 26(9): 191-206.

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. (Ed.). *Supercentenarians (Demographic Research Monographs 7)* (pp. 219-230). Berlin: Springer.

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. & Coale, A. J. (1990). Age patterns of mortality for older women: An analysis using the age-specific rate of mortality change with age. *Mathematical Population Studies*, 2(4): 245-267.

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. (1992). Frailty and survival. *Genus*, 47(3-4): 101-118.

Kannisto, V. (2000). Measuring the compression of mortality. *Demographic Research,* 3(6), doi: 10.4054/DemRes.2000.3.5.

Kestenbaum, B. (1992). A description of the extreme aged population based on improved Medicare enrollment data. *Demography,* 29**,** 565-580.

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.

Masters, R. K. (2012). Uncrossing the U.S. black-white mortality crossover: The role of cohort forces in life course mortality risk. *Demography*, 49(3): 773-796.

Missov, T. I. & Finkelstein, M. S. (2011). Admissible mixing distributions for a general class of mixture survival models with known asymptotics. *Theoretical Population Biology*, 80(1): 64-70.

Mueller, L. D., Rauser, C. L., & Rose, M. R. (2011). *Does aging stop?* Oxford: Oxford University Press.

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

Palloni, A. (2006). Reproducing inequalities: Luck, wallets, and the enduring effects of childhood health. *Demography,* 43**,** 587-615.

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?* Paper presented at the Population Association of America. Detroit.

Rauser, C. L., Abdel-Aal, Y., Sheih, J. A., Suen, C. W., Mueller, L. D., & Rose, M. R. (2005). Lifelong heterogeneity in fecundity is insufficient to explain late-life fecundity plateaus in Drosophila melanogaster. *Experimental Gerontology*, 40 (8-9): 660-670.

Robert, S. A. & House, J. S. (2000). Socioeconomic inequalities in health: An enduring sociological problem. In Bird, C. E., Conrad, P. & Fremont, A. M. (Ed.), *Handbook of Medical Sociology, 5th edition* (pp. 79-97). Upper Saddle River: Prentice Hall.

Rodriguez, G. (1994). Statistical issues in the analysis of reproductive histories using hazard models. *Annals of the New York Academy of Sciences*, 709: 266-279.

Rogers, R. G. (2002). Mortality differentials in a diverse society. In Denton, N. A. & Tolnay, S. (Ed.), *American Diversity: A demographic challenge for the twenty-first century* (pp. 129-154). Albany: SUNY Press.

Steinsaltz, D. (2005). Re-evaluating a test of the heterogeneity explanation for mortality plateaus. *Experimental Gerontology*, 40(1-2): 101-113.

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

Steinsaltz, D. R., & Wachter, K. W. (2006). Understanding mortality rate deceleration and heterogeneity. *Mathematical Population Studies*, 13(1): 19-37.

Strehler, B. L., and Mildvan, A. S. (1960). General theory of mortality and aging. *Science,* 132(3428): 14-21.

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. (Ed.) *Between Zeus and the salmon: The biodemography of longevity* (pp. 17-37). Washington, D.C.: National Academy Press.

Vaupel, J. W. and J. R. Carey. (1993). Compositional interpretations of medfly mortality. *Science*, 260: 1666-1667.

Vaupel, J. W. & Gampe, J. (2009). *Supercentenarians and the theory of heterogeneity*. Paper presented at the International Population Conference. 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.

Wachter, K. W. & Finch, C. E. (Ed.) (1997). *Between Zeus and the salmon: The biodemography of longevity.* Washington, D.C.: National Academy Press.

Zajacova, A., Goldman, N., & Rodriguez, G. (2009). Unobserved heterogeneity can confound the effect of education on mortality. *Mathematical Population Studies*, 16(2): 153-173.

Zheng, H., Yang, Y., & Land, K. C. (2011). Heterogeneity in the Strehler-Mildvan general theory of mortality and aging. *Demography*, 48: 267-90.

Tables

Table 1. Three definitions of deceleration relative to Gompertz mortality

	Mortality, $\mu(a)$		
	Slope, $\mu'(a)$	Acceleration, $\mu''(a)$	Jerk, $\mu'''(a)$
Increasingly accelerating mortality	>0	>0	>0
Relative deceleration	>0	>0	**<0**
Absolute deceleration	>0	**<0**	
Mortality decline	**<0**		

Table 2. Proportion of cohorts for which mortality and age at deceleration increase with baseline percent frail

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

Figures

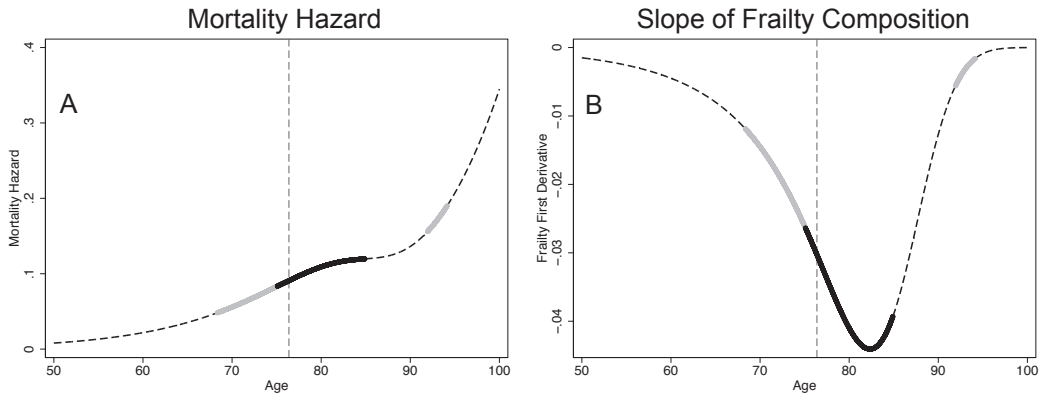

Figure 1. Deceleration Intervals in Example Cohort. The left column gives the mortality hazard and the right column gives the slope of the frailty composition (the negative rate of mortality selection), both over age. The dashed dark lines represent Gompertz mortality; the thick grey lines, absolute deceleration; and the thick black lines, relative deceleration. The dashed light vertical line marks the point where the frail become a minority.

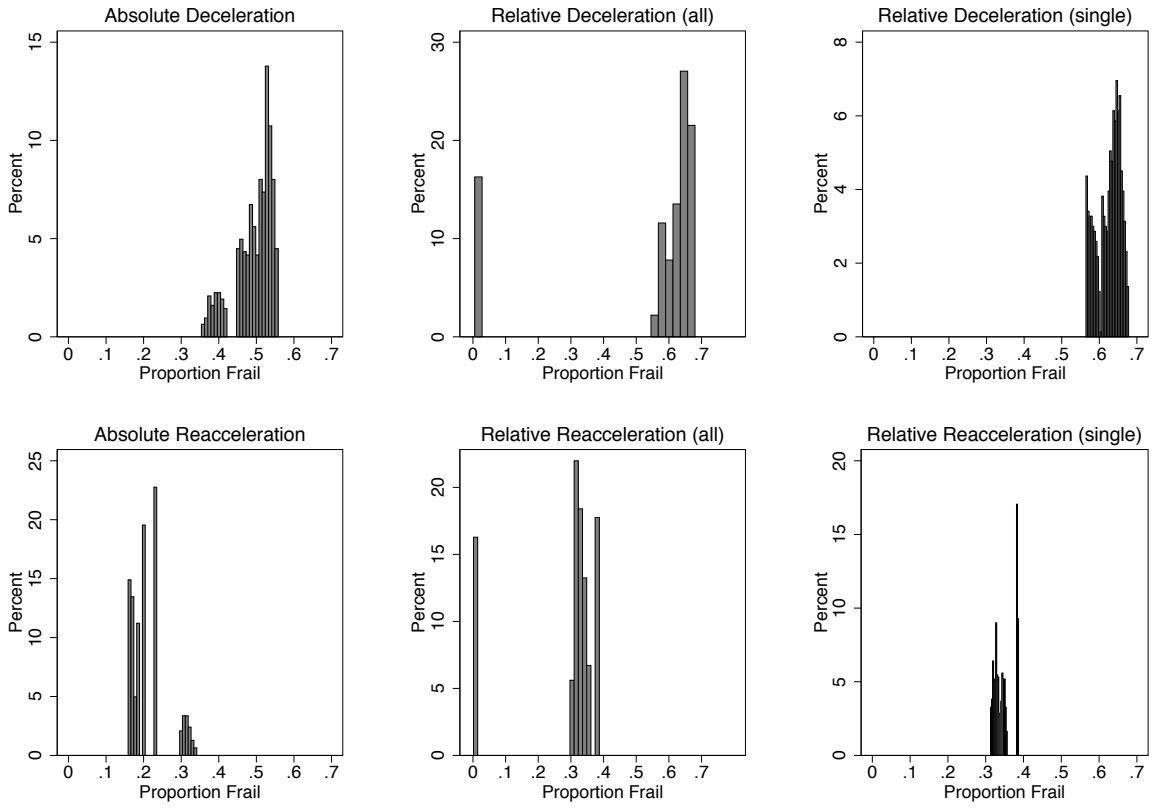

Figure 2. Proportion Frail at Deceleration/Reacceleration. The top row is deceleration, and the bottom reacceleration; the columns are, respectively, absolute deceleration, all relative decelerations, and relative decelerations limited to cohorts that decelerate only once.

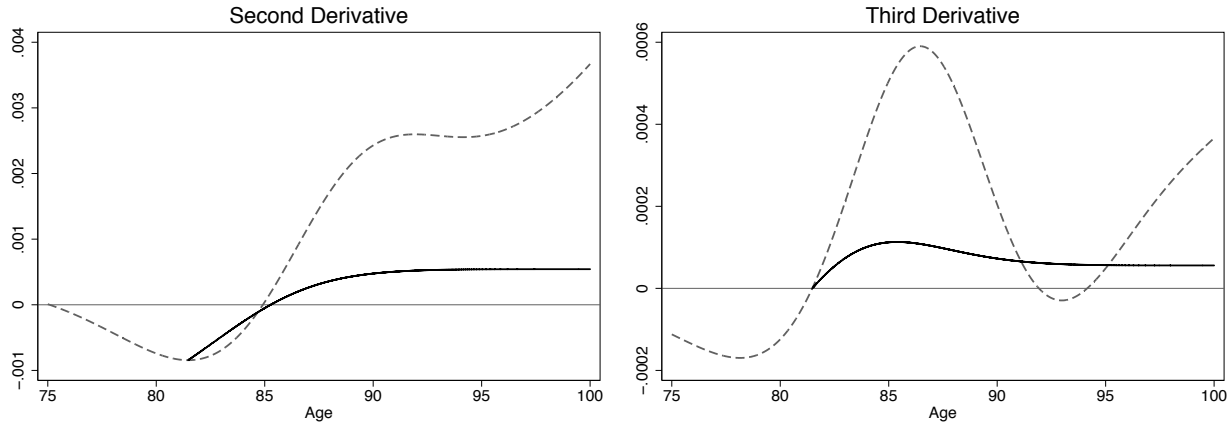

Figure 3. Acceleration Caused by Mortality Selection. The solid black line gives the artificial derivatives calculated by fixing subpopulation mortality and allowing frailty composition to

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

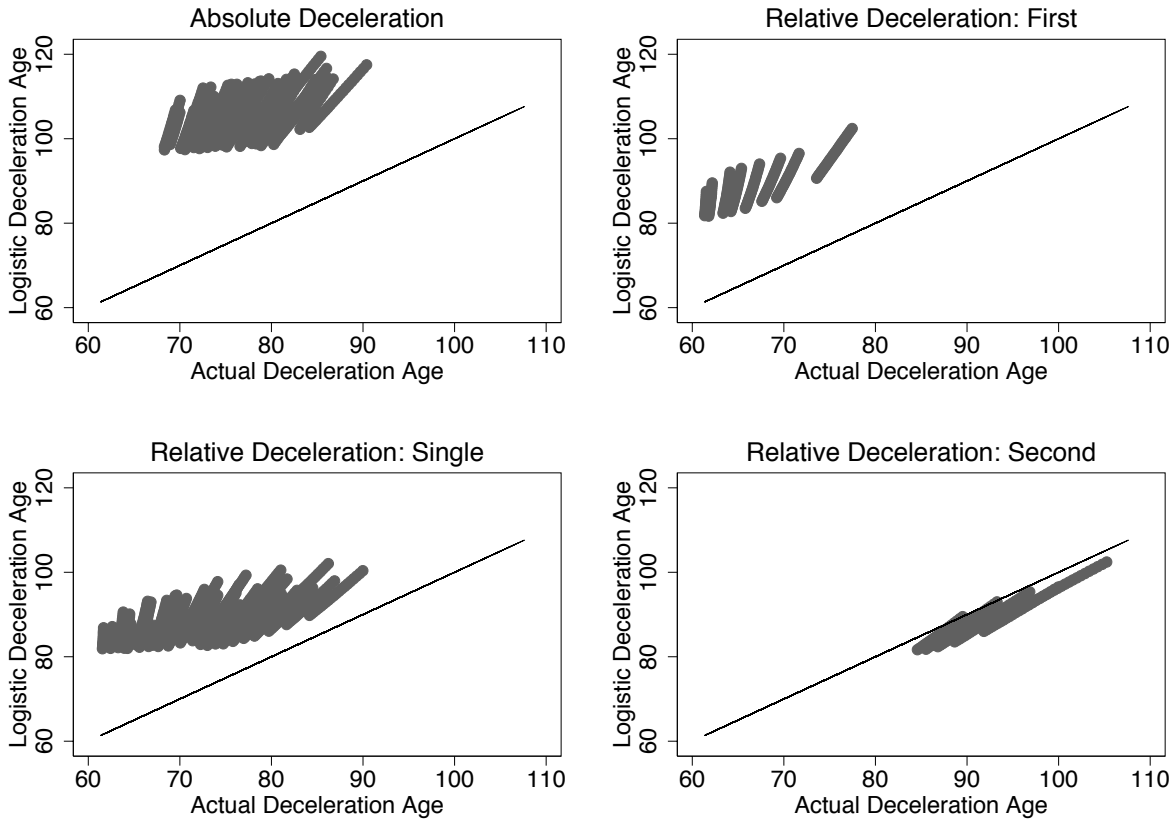

Figure 4. Deceleration Age: Predicted from Logistic Models vs. Actual. The main diagonal is drawn for reference.

Supplementary Appendix

Appendix 1

Here I present the equations that fully describe population deceleration for the model given in the paper. To illustrate, I also visually present the first, second, and third derivatives of the aggregate mortality hazard and the frailty composition for the example cohort discussed in the paper (defined by α=.002, β=.103, f=5, π_0=.75; the derivatives are shown in Fig. 1 and summarized in Table 1, and discussed below).

The model defines frail mortality as proportional to robust mortality, both increasing exponentially over age (i.e., both subpopulations have Gompertz mortality), as shown in Equations 1a and 1b:

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

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

Cohort mortality at any age is the average mortality of the two subpopulations, weighted by their proportion in the cohort, as shown in Eq. 2:

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

In this and all subsequent equations, the first line presents the expression as a function of frailty composition, $\pi(x)$; frail and robust subpopulation mortality, μ_f and μ_r; and (later) their respective derivatives, while the second line presents the expression as an equivalent function of frailty composition and its slope, robust mortality, and the two key subpopulation parameters: the frailty multiplier *f*, and (later) the log-slope of the subpopulations, *β*. The two forms allow different insights about when mortality may be expected to accelerate or decelerate. Panel A of Fig. 1 shows this aggregate hazard for the example cohort, discussed in the main text.

The slope of cohort mortality is a function of the slope of the mortality of each subpopulation, the difference between frail and robust mortality, and the slope of the frailty composition (the percent frail), as shown in Eq. 3:

$$\begin{aligned}\overline{\mu}'(x) &= \pi'(x)\left(\mu_f(x)-\mu_r(x)\right)+\pi(x)\mu_f'(x)+\left(1-\pi(x)\right)\mu_r'(x)\\ &= \mu_r(x)\left[\left(f-1\right)\left(\pi'(x)+\beta\pi(x)\right)+\beta\right]\end{aligned} \tag{3}$$

Considering the first line of Eq. 3, readers may recognize this expression for $\overline{\mu}'(x)$ 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 ($\pi(x)\mu_f'(x)+\left(1-\pi(x)\right)\mu_r'(x)$) 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 B of Fig. 1 displays the slope of mortality for the

aggregate cohort; its dynamics will become more interpretable as we analyze the slope of frailty composition.

The slope of the frailty composition is itself a function of the difference between frail and robust mortality, and of the level of the frailty composition, as given in Eq. 4 (the first line of which is identical to Eq. 3 in the main text):

$$\begin{aligned}\pi'(x) &= -\pi(x)\big(1-\pi(x)\big)\Big(\mu_f(x)-\mu_r(x)\Big) \\ &= \Big(\pi(x)^2-\pi(x)\Big)\big(f-1\big)\mu_r(x)\end{aligned} \tag{4}$$

As described in the main text, the relationship between these terms in the slope of the frailty composition—or its absolute value, what I call the *rate of frailty decline*, or the *rate of mortality selection*—provides the intuition for the simulation results. Panel E of Fig. 1 displays the frailty composition for the example cohort, and Panel F displays its slope.

The second derivative of the frailty composition describes whether the rate of mortality selection is increasing (when $\pi''(x)<0$, since the rate of selection is $-\pi'(x)$) or decreasing ($\pi''(x)>0$). Equation 5 gives this expression:

$$\begin{aligned}\pi''(x) &= \Big(\pi^2(x)-\pi(x)\Big)\Big(\mu_f'(x)-\mu_r'(x)\Big) \\ &\quad +\pi'(x)\big(2\pi(x)-1\big)\Big(\mu_f(x)-\mu_r(x)\Big) \\ &= \pi'(x)\Big[\big(2\pi(x)-1\big)\big(f-1\big)\mu_r(x)+\beta\Big]\end{aligned} \tag{5}$$

In the first form of Eq. 5, the first term, $\left(\pi^2(x)-\pi(x)\right)\left(\mu_f'(x)-\mu_r'(x)\right)$, 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 second form of Eq. 5 can be rearranged to show more specifically that $\pi''(x)>0$—that is, the rate of mortality selection decreases over age—when $\pi(x)<\frac{1}{2}-\frac{\beta}{2\left(\mu_f(x)-\mu_r(x)\right)}$. The greater the ratio of the log-slope of mortality for each subpopulation to the difference between frail and robust mortality at the given age, the farther the frailty composition must fall below half of the cohort for the rate of selection to slow down. In the example cohort, as shown in Panel G of Fig. 1 and in the upper right panel of Table 1, the rate of selection declines after age 82, at frailty composition .27.

The third derivative of the frailty composition describes whether the rate of frailty decline is accelerating ($\pi'''(x)<0$) or decelerating ($\pi'''(x)>0$). Equation 6 gives this expression:

$$\begin{aligned}\pi'''(x) &= \pi''(x)\left(2\pi(x)-1\right)\left(\mu_f(x)-\mu_r(x)\right)+2\left(\pi'(x)\right)^2\left(\mu_f(x)-\mu_r(x)\right)\\ &\quad +2\pi'(x)\left(2\pi(x)-1\right)\left(\mu_f'(x)-\mu_r'(x)\right)+\left(\pi(x)^2-\pi(x)\right)\left(\mu_f''(x)-\mu_r''(x)\right)\\ &= \left(f-1\right)\mu_r(x)\begin{bmatrix}2\left[\left(f-1\right)\mu_r(x)\right]^2\left[\pi(x)^4-2\pi(x)^3+3\pi(x)^2-2\pi(x)+1\right]\\ +3\beta\left(f-1\right)\mu_r(x)\left(2\pi(x)-1\right)\left(\pi(x)^2-\pi(x)\right)\\ +\beta^2\left(\pi(x)^2-\pi(x)\right)\end{bmatrix}\end{aligned} \quad (6)$$

The sign of the third derivative of the frailty composition is the sign of the main bracketed term in the second form of this expression. Within these brackets, the first term, $2\left[\left(f-1\right)\mu_r(x)\right]^2\left[\pi(x)^4-2\pi(x)^3+3\pi(x)^2-2\pi(x)+1\right]$, is always positive, the third term, $\beta^2\left(\pi(x)^2-\pi(x)\right)$, is always negative, and the middle term, $3\beta\left(f-1\right)\mu_r(x)\left(2\pi(x)-1\right)\left(\pi(x)^2-\pi(x)\right)$, is negative when $\pi(x)>.5$ and positive when $\pi(x)<.5$. 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 the example cohort, as shown in Panel H of Fig. 1 and the lower right panel of Table 1, the third derivative of frailty composition switches sign just after $\pi(x)=.5$: it is negative (the rate of selection accelerates) until age 77, when the frail are 47% of the cohort. It then remains positive (the rate of selection decelerates) until age 87, when only 7% of the cohort is frail, after which point in remains negative but approaches zero as the frail become extinct.

Returning to cohort mortality, Table 1 and the bottom four panels of Fig. 1 highlight that the dynamics of the second and third derivatives of cohort mortality (shown on the left), whose signs respectively define absolute and relative deceleration, are heavily driven by the second and third derivatives of frailty composition (shown on the right).

Mortality decelerates absolutely when the second derivative of cohort mortality is negative. Equation 7 gives this second derivative of cohort mortality with respect to age:

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

The third term of the first form of the expression in Eq. 7, $\pi(x)\mu_f''(x)+\left(1-\pi(x)\right)\mu_r''(x)$, representing the composition-weighted increase in subpopulation slopes, 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 C of Fig. 1 and in the upper left panel of Table 1, mortality decelerates relatively at age 75, when the frail are 54% of the cohort, and reaccelerates at age 84, when the frail are 16% of the cohort.

Mortality decelerates relatively when the third derivative of cohort mortality is negative. Equation 8 gives the third derivative of cohort mortality with respect to age:

$$\begin{aligned}\overline{\mu}'''(x) &= \pi'''(x)\left(\mu_f(x)-\mu_r(x)\right)+3\pi''(x)\left(\mu_f'(x)-\mu_r'(x)\right)\\ &\quad +3\pi'(x)\left(\mu_f''(x)-\mu_r''(x)\right)+\pi(x)\mu_f'''(x)+\left(1-\pi(x)\right)\mu_r'''(x)\\ &= 2\left[\left(f-1\right)\mu_r(x)\right]^4\left[\left(2\pi(x)-1\right)^2+2\left(\pi(x)^2-\pi(x)\right)^2\right]\\ &\quad +6\beta\left[\left(f-1\right)\mu_r(x)\right]^3\left(2\pi(x)-1\right)\left(\pi(x)^2-\pi(x)\right)\\ &\quad +9\beta^2\left[\left(f-1\right)\mu_r(x)\right]^2\left(\pi(x)^2-\pi(x)\right)\\ &\quad +\beta^3\left(f-1\right)\mu_r(x)\cdot\pi(x)+\mu_r(x)\end{aligned} \tag{8}$$

In the first form of this expression, the fourth term, $\pi(x)\mu_f'''(x)+\left(1-\pi(x)\right)\mu_r'''(x)$, representing the composition-weighted increase in subpopulation acceleration, is always positive, and the 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 change in frailty, 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 D of Fig. 1 and in the lower left panel of Table 1, mortality decelerates relatively at age 68, when the frail are 66% of the cohort; reaccelerates at age 81, when the frail are 31%; decelerates relatively a second time at age 91, when the frail are only 1%; and reaccelerates a final time at age 94, when the frail are only two-tenths of 1% 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 mortality selection 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 ultimate patterns may depend heavily on the values of the subpopulation mortality parameters.

References

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

Table

Table 1. Turning points 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	+			

Figure

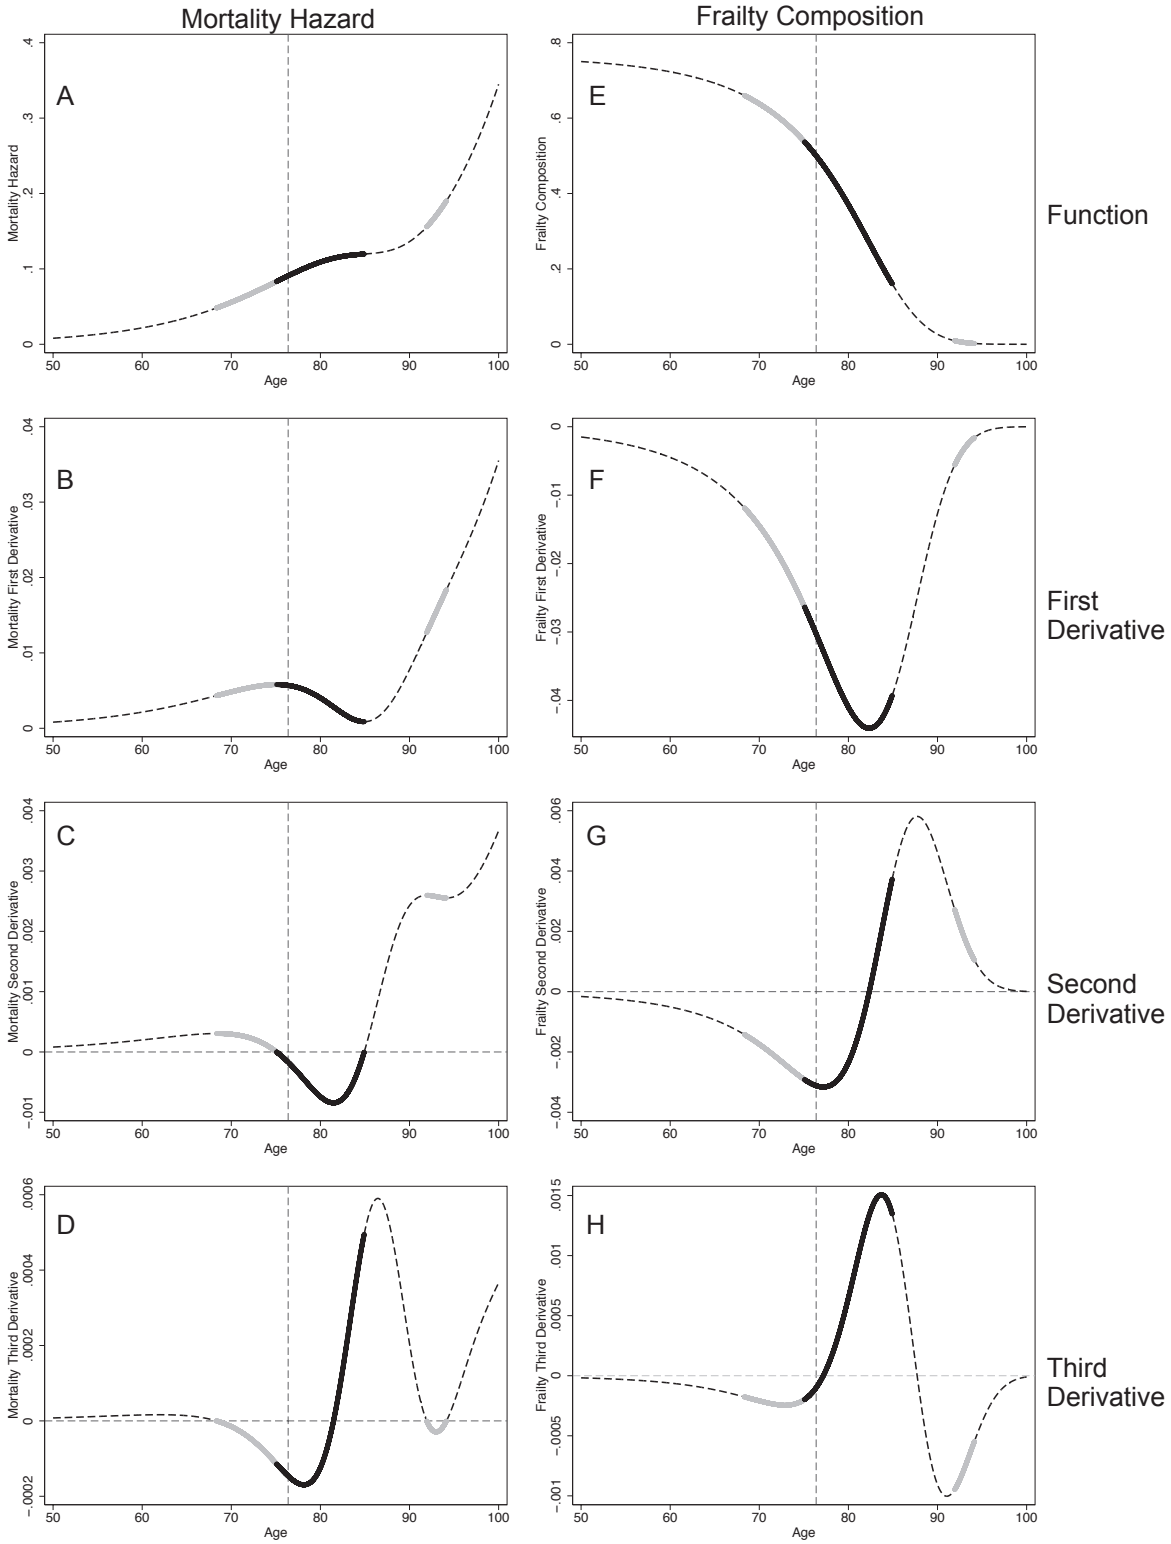
Mortality Hazard
Frailty Composition
A
E
B
F
C
G
D
H
Mortality Hazard
Frailty Composition
Mortality First Derivative
Frailty First Derivative
Mortality Second Derivative
Frailty Second Derivative
Mortality Third Derivative
Frailty Third Derivative
Age
Function
First Derivative
Second Derivative
Third Derivative

Figure 1. Deceleration Intervals in Example Cohort. The left column gives mortality and the right column gives frailty composition (proportion frail), 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.

Appendix 2

Here I analyze the parameter space of high-frailty and multiple mortality deceleration, in terms of the intercept, log-slope, and frailty multiplier of the subcohorts.

Fig. 1, Panel A, displays in grey the aggregate Gompertz parameters of the universe of 1,151 simulated cohorts analyzed in this paper, with those of the HMD cohorts overlaid in black.[20] This panel illustrates visually the tight fit between the aggregate parameters of the real and simulated cohorts. Among the simulated cohorts, the robust subcohort intercepts range from .001 to .016, and the subcohort log-slopes from .062 to .113. At each aggregate Gompertz intercept value, the range of admissible aggregate Gompertz log-slopes has a spread of .02.

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 1. 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 1, 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 aggregate mortality values ranging from .07 to .16. As shown in the second column of Table 1, *all* of the 910 cohorts in which the frail subpopulation has at least 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

[20] 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 (Engelman et al. 2010, Kannisto 2000).

mortality values ranging from .04 to .20. Mortality decline is not found in the universe of simulated cohorts close to the HMD cohorts.

The incidence of high-frailty deceleration, which I define for this analysis as deceleration when the frail are a majority of the cohort, is shown in Panels B and C of Fig. 1. Panel B shows simulated cohorts evincing high-frailty *absolute* deceleration. Such decelerations are found in 30% of total cohorts in this universe, and 56% 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 1, *all* cohorts with frailty multiplier equal to at least 4 evince high-frailty absolute deceleration.

Panel C of Fig. 1 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 1, 100% 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.

The incidence of multiple deceleration—which occurs only for relative, not absolute, deceleration—is shown in Panel D of Fig. 1. 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—15% 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.

References

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.

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

Table

Table 1. 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

Figure

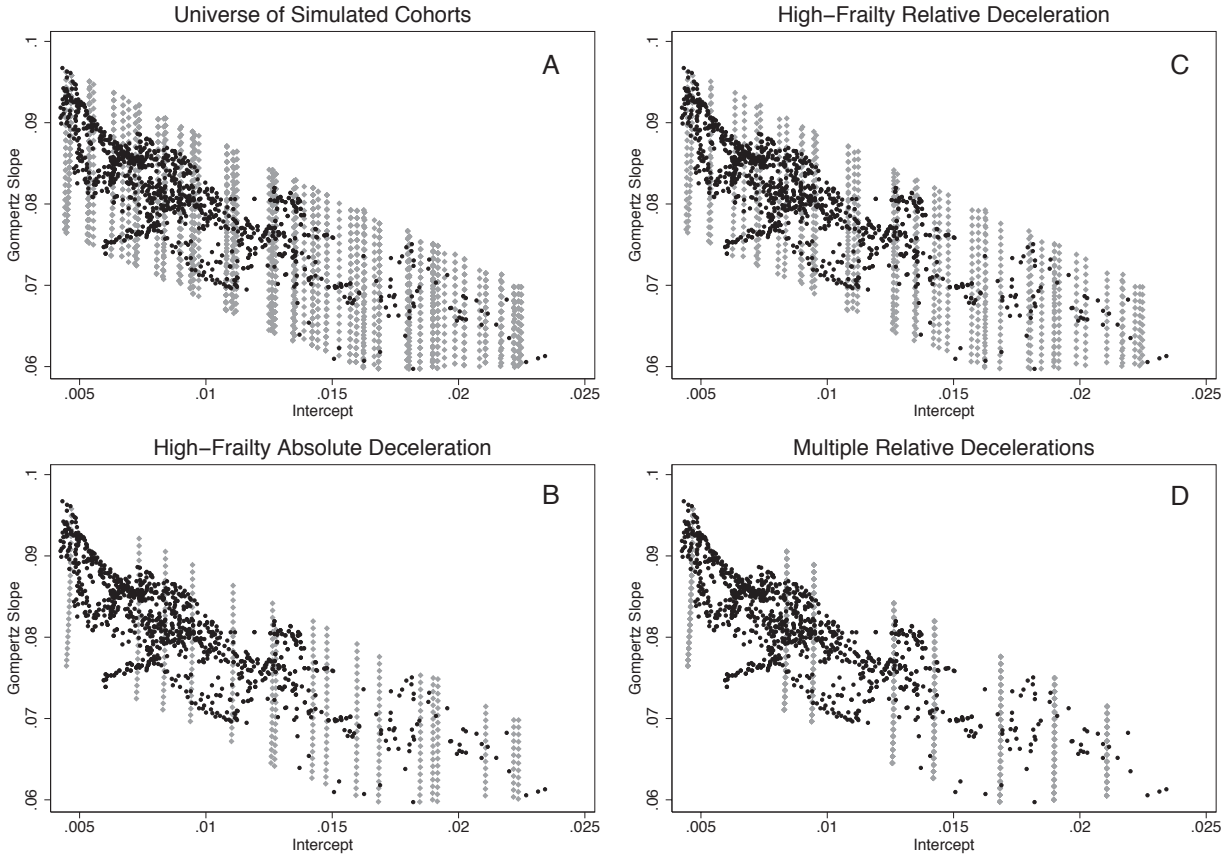

Figure 1. Gompertz Parameters of Simulated and Real Cohorts. Simulated cohorts are shown in grey and cohorts from the Human Mortality Database are shown in black. (A) Universe of all simulated cohorts. (B) Simulated cohorts with high-frailty absolute deceleration (absolute deceleration while most of the cohort is frail). (C) Simulated cohorts with high-frailty relative deceleration. (D) Simulated cohorts with multiple relative decelerations.

Appendix 3

Here I present an example of the cross-cohort comparisons discussed in the paper, and add as a dimension of comparison—in addition to the age and aggregate mortality at deceleration—the percent frail at deceleration.

To generate the example set of cohorts, 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.[21] This example is one set of α, β, and f combinations among the 1,151 combinations (representing the full universe of HMD-compatible simulations) used to generate the cross-cohort comparisons reported in the main paper.

Fig. 1 displays the results. The panel 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 panel 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.

Four important results are suggested in Fig. 1 (and confirmed in the global analysis summarized in the main text). 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

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

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 decelerations are distinguished from one another.[22] 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 (in data truncated by age on either or both ends) 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.

In contrast to these results for mortality and age, all four of these types of deceleration—absolute, single relative, 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. This result is not an artifact of the choice of example set of cohorts: in the full set of cohort comparisons, the positive association between the baseline

[22] 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 of 99%.

frailty composition and the frailty composition at deceleration *always* holds, as long as first and second relative decelerations are distinguished from one another.

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 these results suggest because real data are truncated at the oldest ages. The analysis given here 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.

References

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.

Figure

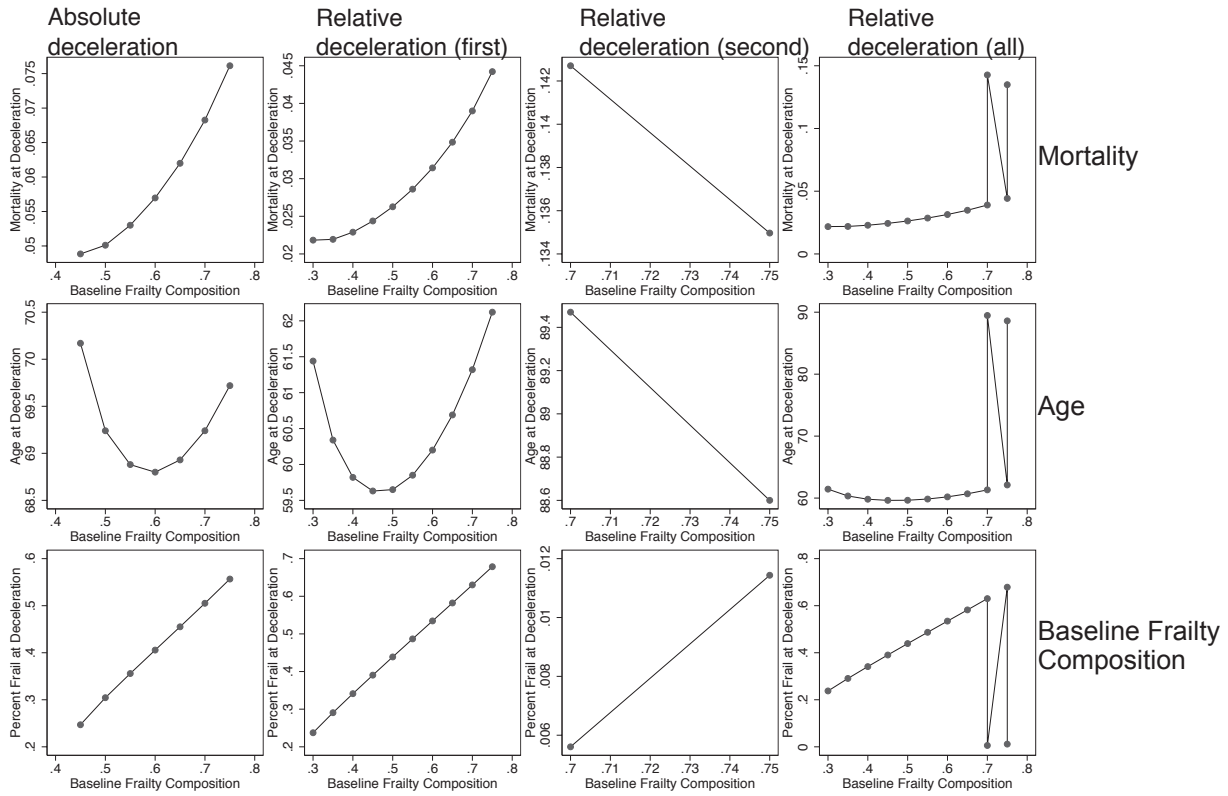

Figure 1. Relationship of Baseline Frailty Composition to Mortality, Age, and Frailty Composition at Deceleration.

CHAPTER 2:

Mortality Deceleration is Not Informative of

Unobserved Heterogeneity in Open Groups

CHAPTER ABSTRACT

Studies of mortality deceleration sometimes use the order of deceleration between two groups—such as birth cohorts or subpopulations defined by race or sex—to draw conclusions about population heterogeneity. These studies often draw on the fact that, between two closed groups, ceteris paribus, the higher-mortality group will decelerate at a younger age. This paper gives a first consideration to the order of deceleration between open groups. I construct a model with a one-way flow from "healthy" to "sick" status and, crucially, assume that a subset of the cohort has elevated risk for both sickness and death. Using simulations designed to resemble, in the aggregate, cohorts in the Human Mortality Database, I show that mortality deceleration order in such a cohort is essentially unpredictable because it depends on the interaction between a large number of parameters, some of which are unobserved in empirical data. These results suggest that it may be challenging to extend the study of mortality deceleration to include open groups whose memberships are selected for frailty or robustness, while still drawing meaningful conclusions about population heterogeneity.

Introduction

Cohort mortality often decelerates: the rate of mortality increase over age slows down at old ages, producing a plateau in some cohort mortality hazards. This deceleration is often attributed, in whole or in part, to mortality selection: the change in cohort composition over age as "frailer" members die, leaving an increasingly robust, low-mortality cohort. Most empirical research on mortality deceleration focuses on decelerations in closed cohorts or closed groups defined by race or sex. Selection dynamics in such closed heterogenous groups are well explored in the theoretical literature (e.g., Beard 1959, 1971; Kannisto 1992; Vaupel, Manton, and Stallard 1979; Thatcher et al. 1998; Vaupel and Yashin 1985), and empirical research profitably draws on these theoretical results.

Chief among the theoretical predictions that guide empirical work on mortality deceleration is that, if mortality deceleration results from mortality selection, then higher mortality groups, ceteris paribus, should decelerate at younger ages than lower mortality groups. This stems from the fact that higher-mortality groups are subject to more intense selective pressure (Vaupel and Yashin 1985; Vaupel, Manton and Stallard 1979; Kannisto 1992). Such predictions about the order of groups' mortality deceleration allow empirical comparisons of deceleration timing to be used in one of two ways.

First, such predictions allow selection-based explanations of deceleration to be tested, an important undertaking because a large body of biodemographic mounts a case for decelerating mortality in individuals, not only populations (e.g., Carey et al. 1992, Curtsinger et al. 1992, Drapeau et al. 2000 [but see Steinsaltz 2005], Mueller et al. 2011, Rauser et al. 2005, Steinsaltz and Wachter 2006, Vaupel and Carey 1993; and see reviews in Vaupel 1997, Wachter and Finch

1997). For example, Horiuchi and Wilmoth (1998) find that the age of deceleration onset has risen over time as the level of cohort mortality has fallen (see also Engelman [2010]), and take this cross-cohort comparison as evidence in favor of the selection explanation, since it is consistent with what that explanation would predict. Similarly, Lynch and Brown (2001) find that, within the United States, higher-mortality male cohorts decelerate at younger ages than lower-mortality women, and consider this evidence for selection.

Second, such predictions allow analysts to assume that differences in deceleration timing arise from differential selection, and therefore use empirical data to derive information about latent heterogeneity. For example, Lynch et al. (2003) find later deceleration among black Americans compared to white Americans in spite of the black cohorts' greater mortality, and conclude that black cohorts born 1972-1990 had greater frailty from birth than white cohorts: selection theory predicts that if the groups had the same frailty distribution at birth, then the black cohorts should decelerate earlier, not later.[1] In sum, the generalization that, all else equal, higher-mortality groups should decelerate at younger ages is a linchpin connecting mortality selection theory to the empirical testing of its assumptions.

This paper explores the selection dynamics of mortality deceleration between open groups, i.e., groups that can be entered as well as exited, such as sick vs. healthy, poor vs. non-poor, or countries that exchange members. Specifically, I ask whether it remains true for open groups that higher mortality groups should decelerate before lower mortality groups.

This question is motivated in part by a growing interest among demographers in expanding the study of mortality to include health and disability statuses (e.g., Robine et al.

[1] As Lynch et al. note, this conclusion rests on two further assumptions as well: that frailty has the same mortality consequences in each group, and that the groups age at the same rate. (These conditions are formalized below, when I present the models.)

2003, Zeng et al. 2006). In part, this reflects the increasing availability of longitudinal data with rich health covariates. It also reflects the concerns of aging societies that, hoping to enable interventions supporting increased quality of life at old ages, may conceptualize mortality not merely as an event, but as a process that may—or may not—include varying dimensions of morbidity before culminating in death. An early tradition of mortality research, centered around Manton, Stallard, Woodbury, and Yashin (e.g., Manton et al. 1994, 1995; Woodbury and Manton 1983), did integrate health transitions with individual heterogeneity in a variety of models, but largely has not been echoed by more recent demographic research.[2] Perhaps because such efforts have not gained wider traction, it remains unknown what generalizations about phenomena such as mortality deceleration might apply in those contexts. As mortality research continues to incorporate health and morbidity, an open question is what kinds of knowledge about mortality, rooted in comparisons of relatively closed groups such as cohorts and countries, can be extended to open groups such as health statuses that individuals move between.

I will address this question on the most favorable ground for mortality selection theory, in order to dramatize the difficulties that exist, even on that ground, in applying selection theories about mortality deceleration to open groups. A large body of biodemographic research (cited above) challenges the idea that deceleration results solely from selection; I will assume that it does. Previous work (Wrigley-Field 2013) also shows that some widespread generalizations about mortality deceleration need not hold even in the closed-group context; the generalization considered here—that higher-mortality groups decelerate first—does. I also use a very simple model in which frailty is binary and fixed in individuals (following the classic early work

[2] More generally, there is a small but important literature on heterogeneity in dynamic settings (particularly Mohtashemi and Levins [2002] and Rogers [1992]) that, while concerned with different issues than mortality deceleration, can broadly be considered inspiration for this paper.

building demographic intuitions about closed groups, Vaupel and Yashin [1985]), and there is only a one-way flow from sickness to health. In short, the analysis in this paper is designed to be favorable to finding straightforward predictions about mortality deceleration in an open-group context. It would be desirable for the predictions about deceleration order among open groups to be similarly straightforward as the closed-group generalization that higher-mortality groups decelerate at younger ages. This paper demonstrates that they are not.

In what follows, I begin by defining mortality deceleration and then by more precisely presenting the closed-group and open-group selection models. I then intuitively motivate a hypothesis: that the order of deceleration is close to unpredictable in the dynamic (open group) context because it is very sensitive to the relative sizes of the subgroups, including latent subgroups. To evaluate this hypothesis, I present a series of simulations designed to resemble, in the aggregate, cohorts in the Human Mortality Database.

Defining deceleration

This paper defines a deceleration point as a point when the second derivative of unlogged mortality becomes negative (or in other words, when the mortality slope begins to decline). This is the empirical measure used in the work (Lynch et al. 2003; see also Lynch and Brown 2001) that, by explicitly stating the conditions (reviewed below) for higher-mortality closed groups to decelerate first, helps to inspire the current analysis.

A more commonly-used measure is the Lifetable Aging Rate (LAR), the slope of logged mortality (introduced by Horiuchi and Coale [1990], though the term is introduced later; see also Horiuchi [1997], Horiuchi and Wilmoth [1997]). By taking the derivatives of logged mortality,

the LAR focuses attention on the increase in mortality relative to the overall level of mortality. This relative measure has the advantage of capturing an intuitive idea of what mortality deceleration means (Horiuchi and Coale 1990), but the disadvantage that selection dynamics are conflated with the level of mortality (Vaupel and Zhang 2010). For this paper, which is concerned with understanding the relationship between selection dynamics and deceleration order, this disadvantage is paramount, and so measures that take the derivatives of mortality directly are preferable to the LAR.[3]

The closed-group and open-group selection models

Previous work on deceleration uses closed-group models in which higher-mortality groups decelerate first. The main purpose of this paper is to examine a health selection model in which some subset of the population has elevated risks both for becoming sick and for mortality conditional on sickness. This new model incorporates two changes, compared to the closed-group model: the addition of a one-way flow from health to sickness, and the fact that the frail are at heightened risk for the new flow (as well as mortality). To isolate the effects of each change, I also consider an alternative open-group model in which the healthy can become sick, but—unlike the main model—the robust and the frail are equally likely to transition from health to sickness. These three models are summarized in Figure 1, and are given formally below.

Substantively, the closed-group model can be thought of as a model in which sickness/health is assigned at birth rather than acquired, and the two open-group models are

[3] There is also an alternative measure using the mortality derivatives: Rau (2009) proposes marking deceleration as the point when the *third* derivative of mortality becomes negative. I avoid this measure in this paper because it admits the possibility of multiple decelerations even in the closed-group case with binary frailty (Wrigley-Field 2013), which would complicate this analysis with little analytical payoff for present purposes.

distinguished by whether the risk factors for sickness and for mortality (conditional on sickness) overlap. The main model, the open-group model with elevated sickness risk for the frail, posits that there are some shared risks for sickness and for death operating at all ages. For example, smoking, living in poverty, or being born low-birthweight might each increase the risk of heart disease, and also increase the risk of death regardless of whether one has heart disease. The comparison model, in which the frail and robust have equal transition rates to sickness, posits that there are no shared risks for sickness and for mortality conditional on sickness. This might be a reasonable approximation for a population in which most deaths are to accidents and the sick are no more likely to engage in risky behavior, but for most purposes, the main model seems a better approximation.

As we will see, just transitioning to an open group creates problems for the predictability of deceleration order, but the major problems come with shared risks for sickness and death.

Closed-group model

Consider two closed groups, say, men and women. Imagine that both groups are heterogeneous, composed of two kinds of people—the "frail" and the "robust"—with frailty fixed at birth. (In a typical empirical context, such frailty would be unobserved and posited theoretically.) All individuals have Gompertz mortality with the same slope over age. The frail have higher mortality than the robust, by the same proportion for men as for women, and men have higher mortality than women. Aggregate mortality in each group depends on the mortality of frail and robust individuals, and on the proportion of each group (men or women) that is frail at each age. The individual-level mortality is given in Equation 1:

$$\mu_{robust\ women}(a) = \alpha e^{\beta a}$$
$$\mu_{frail\ women}(a) = f\alpha e^{\beta a}$$
$$\mu_{robust\ men}(a) = m\alpha e^{\beta a} \quad (1)$$
$$\mu_{frail\ men}(a) = fm\alpha e^{\beta a}$$

where β is the log-slope of mortality for all individuals, α is the intercept of mortality for robust women, f is the frailty multiplier on mortality, and m is the male multiplier on mortality. If the proportion frail at any age is denoted $\pi(a)$, then the mortality over age for women and men is given by Equation 2:

$$\bar{\mu}_{women}(a) = (1-\pi_{women}(a))\alpha e^{\beta a} + \pi_{women}(a) f\alpha e^{\beta a}$$
$$\bar{\mu}_{men}(a) = (1-\pi_{men}(a))m\alpha e^{\beta a} + \pi_{men}(a) mf\alpha e^{\beta a}, \text{ for } f, m > 1 \quad (2)$$

The percent frail in each group, , $\pi_{women}(a)$ and $\pi_{men}(a)$, can be calculated directly from the survivorships at each age for each subgroup defined by sex and frailty. The survivorships for each subgroup *i* at age *a*+*x* are given in terms of the survivorships at age a in Equation 3:

$$S_i(a+x) = S_i(a) \cdot \left(1 - \mu_i(a)x\right) \quad (3)$$

In this closed model, mortality is the only source of change in the size of each subgroup, and hence of the frail/robust composition of the groups.

The generalization that men will decelerate at a younger age than women rests on three ceteris paribus assumptions built into this model (given in Lynch et al. [2003]): the groups share an individual-level log-slope of mortality over age, β; they share the same proportional inequality between frail and robust mortality, f; and, not shown in Equation 1, they must also have the same proportion frail at baseline, $\pi(0)$.[4]

Open-group model

The open-group models considered here retain the three assumptions from Lynch et al., but the move from closed to open groups adds substantial complexity to this model. To limit this complexity while honing in on the key features of open groups relevant to mortality selection, the model considered here incorporates several constraints. The most important such constraint is that the open-group model considered here has only a one-way flow from health to sickness. Thus, at any age, the healthy can stay healthy, become sick, or die, whereas the sick can only stay sick or die. A second key assumption of the model is that the frail are more likely than the robust to become sick, as well as to die regardless of whether they are sick. This assumption can be conceptualized as the idea that there are some shared risks for morbidity and, conditional on morbidity, for mortality, that are fixed at the individual level. (Results from models omitting this assumption are also presented below, for comparison.) Finally, the open-group model retains the ceteris paribus conditions for the higher-mortality group to decelerate first in the closed-group

[4] The assumption that the healthy and sick groups have equal frailty composition at baseline is likely not substantively realistic. I nevertheless make it, for two reasons. First, the assumption retains continuity with the closed-group model, which is essential since the question addressed here is whether a generalization about the closed-group model can extend to open groups. Second, it is not clear *a priori* which group should be expected to have a larger number of frail members, if baseline is not birth—precisely because of the complexity of the differential selection of the frail in and out of the groups, as explained in this section.

model: individuals, regardless of subgroup (defined by sickness/health and frailty/robustness), share a log-slope of mortality over age, β; both groups (defined by sickness/health) have the same proportional inequality between frail and robust mortality, f; and both sickness/health groups have the same proportion frail at baseline, $\pi(0)$.

Thus, the mortality functions in this open-group model, conditional on π(a), are the exact equivalent of the mortality function in the closed-group model, with m ('morbidity') now representing the mortality multiplier associated with being sick (the variable name is retained from Equations 1 and 2 since its role in Equation 4 is identical). These open-group mortality functions are given in Equation 4:

$$\begin{aligned} \bar{\mu}_{healthy}(a) &= (1-\pi_h(a))\alpha_\mu e^{\beta_\mu a} + \pi_h(a) f_\mu \alpha_\mu e^{\beta_\mu a} \\ \bar{\mu}_{sick}(a) &= (1-\pi_s(a))m\alpha_\mu e^{\beta_\mu a} + \pi_s(a) m f_\mu \alpha_\mu e^{\beta_\mu a}, \text{ for } f_\mu, m > 1 \end{aligned} \tag{4}$$

However, the survivorship functions, which determine the value of $\pi(a)$ for each group, and therefore weight the subgroups in the mortality functions just presented, are more complex in the open-group model than the closed-group model—especially for the sick. Both survivorship functions depend on the rate at which the healthy become sick, given in Equation 5:

$$\bar{\omega}(a) = (1-\pi_h(a))\alpha_\omega e^{\beta_\omega a} + \pi_h(a) f_\omega \alpha_\omega e^{\beta_\omega a}, \text{ for } f_\omega > 1 \tag{5}$$

Thus, the total rate of becoming sick is an aggregate rate for the frail and robust healthy, with the frail healthy at greater risk in proportion to the frailty multiplier on sickness, f_ω (which need not

equal the frailty multiplier on mortality, f_{μ}). In addition to the key analytical assumptions outlined above, Equation 4 makes one additional assumption for tractability. The rate of becoming sick among the frail and the robust has an equal log-slope over age, $\beta\omega$. This log-slope may be different from the log-slope of mortality over age, $\beta\mu$, but neither can vary between the frail and the robust. The choice of parametric form for the subgroup rate of becoming sick is relatively arbitrary, but here is assumed—like mortality—to be Gompertz to reflect substantively plausible increasing morbidity risk over age. (The open-group comparison model in which the frail and robust are equally likely to become sick is identical to the model given in Equations 4 and 5 except that the frailty multiplier on sickness, $f\omega$, is constrained to equal, not exceed, 1.)

The survivorship of the healthy and sick can then be given in terms of the rate of becoming sick, $\omega(a)$, given just above in Equation 4. The survivorship functions at age $a+x$ for the healthy and sick, for subgroup z (frail/robust), are given in Equation 6:

$$\begin{aligned} S_{h,z}(a+x) &= S_{h,z}(a)\cdot\left(1-\mu_{h,z}(a)x\right)\left(1-\omega_{z}(a)x\right) \\ S_{s,z}(a+x) &= S_{s,z}(a)\cdot\left(1-\mu_{s,z}(a)x\right) \\ &\quad +S_{h,z}(a)\cdot\eta(0)\cdot\left(1-\mu_{h,z}(a)x\right)\cdot\omega_{z}(a)x, \\ \text{for } \eta(0) &= \frac{l_{h,z}(0)}{l_{s,z}(0)} \end{aligned} \tag{6}$$

Equation 6 shows that the survivorship of the healthy at any age is the healthy group at the previous age, minus the decrements to death and to sickness. In contrast, the survivorship of the sick is the sick group at the previous age minus the decrement to sickness, plus the portion of

the healthy group at the previous age that survives and becomes sick. Crucially, the composition of the sick group now depends not only on the rate of flow of the healthy group into the sick group, but also on the size of the healthy group relative to the sick. This relative size, in turn, is a function of the relative sizes of the groups at baseline, $\eta(0)$ (defined here in terms of the standard life table function l), and the survivorship of the healthy group. This dependence of the selection dynamics within the sick on the relative size of the healthy will be a crucial factor in the analysis in this paper. (For the ease of interpretation, while I have presented Equation 6 in terms of the ratio of the size of the healthy group to the size of the sick group, in subsequent discussion and analysis I will use the baseline percent of the total cohort that is healthy rather than this ratio.)

In summary, Equations 5 and 6 make it clear that, in spite of the constraints imposed on the open-group model used here, the mortality of the healthy and—especially—the sick, aggregated over frailty, depends on many more parameters than did the aggregate mortality of closed-group men and women. In place of that closed-group model's five parameters—the intercept (baseline mortality) for robust women, α; the frailty, f, and male, m, mortality multipliers; the log-slope of mortality for all subgroups, β; and the baseline percent frail for both groups, $\pi(0)$—are nine parameters in the open-group model. These nine parameters are: the five parameters analogous to the closed-group model—the intercept (baseline mortality) for the robust healthy, α_μ; the frailty, f_μ, and sick, m, mortality multipliers; the log-slope of mortality for all subgroups, β_μ; and the baseline percent frail for both groups, $\pi(0)$—as well as three additional sickness parameters—the intercept (baseline rate) of becoming sick for the robust, α_ω; the log-slope of becoming sick for the frail and robust, β_ω; and the frailty multiplier on becoming sick, f_ω—and an additional baseline cohort composition parameter: the baseline size of the healthy group relative to the sick,

$\eta(0)$. In short, the addition of a single unidimensional path from healthy to sick adds substantial complexity to the aggregate mortality of each group.

Intuition and hypotheses

In the closed-group model, the higher-mortality men should always decelerate at younger ages than the lower-mortality women when the two otherwise share all mortality parameters. In the open-group model, the situation is more complex. To understand the model, we can ask two questions: how do the mortality of the healthy and sick, in the open-group model, compare to women and men in the closed-group model? And, of paramount importance to this paper, how do the mortality of the healthy and sick compare to one another?

For the healthy, the open-group model is a straightforward extension of the closed-group model, with mortality simply extended to two separate decrements, loss to death and loss to sickness. Thus, the healthy in the open-group model certainly are more selected than are the women in the closed-group model with identical mortality parameters. But are they more or less selected than the open-group sick? Although the sick have higher mortality than the healthy, the two decrements from the healthy group might—or might not—give it a higher total decrement rate than the sick group has to mortality alone.[5] Since both death and sickness occur selectively, more often to the frail, this might in principle lead the healthy group to be more quickly selected than the sick group, and thus to decelerate earlier—the opposite of the pattern in the closed-group model, but arising for the same reason as the pattern in that model.

[5] Moreover, the total decrement rate for the healthy and for the sick might cross over age if one has a larger intercept and the other, a larger slope.

For the sick, however, the increment from the healthy group creates substantial complications. To see why intuitively, think of the sick group at any age as divided along a further dimension: the newly sick and the longstanding sick (a stylized distinction serving as a heuristic for the real, continuously-varying duration of sickness). The longstanding sick may have been disproportionately frail early on, since the frail tend to become sick quickly. But subject to the elevated mortality of sickness, that longstanding sick subgroup may have been whittled down to a smaller, largely robust set of survivors.

The newly sick, on the other hand, have not yet faced the intense selective pressures of sickness for very long. Therefore, one possibility is that, as the frail continue to become sick disproportionately, the continued flow from health into sickness produces a frailty replenishment, as the frail continuously become sick, die, and are replaced in the sick group. Such frailty replenishment, by offsetting the increased mortality selection of the sick, could also lead the sick to decelerate at older ages than the healthy, reversing the closed-group deceleration order.

However, the opposite order of deceleration is also possible. The newly sick, rather than replenishing the frailty of the total sick, can also tilt the composition of the total sick farther toward robustness. The easiest way to see this is to imagine that the healthy group simply runs out of frail members.[6] Then the flow of people from healthy to sick will compound the effects of mortality selection among the sick, potentially precipitating a mortality deceleration. One possible outcome is that, if this fall in frailty composition among the sick occurs quickly enough, it might precipitate a mortality deceleration even if the healthy group did not decelerate while it

[6] If the total decrement of the healthy group is greater than the mortality decrement in the sick group, this can happen while the sick group still has frail members—and hence can still decelerate—despite the equal proportion frail in the groups at baseline.

lost its frail members. In that case, the sick might decelerate while the healthy do not. Alternatively, if we now imagine that the healthy have not fully lost their frail members, but have only lost enough to abruptly lower the percent frail among the sick, then we can imagine that the sick decelerate while the healthy will decelerate later, as they lose still more frail. In that case, the open-group model produces the same deceleration order as the closed-group model, rather than reversing it.

In summary, moving to an open-group model (with a one-way flow from healthy to sick and an elevated transition rate among the frail) should lead the healthy to decelerate earlier than the analogously low-mortality women do in the closed-group model with identical mortality parameters. The sick, on the other hand, might decelerate either earlier or later, in this open-group model, than the comparably high-mortality men in the closed-group model with identical mortality parameters but no increment. Most importantly for our purposes, mortality deceleration among the sick can occur either before or after mortality deceleration among the healthy in the open-group model.

Two additional observations can be made before turning to simulated data. First, the degree to which the increment of the newly sick changes the composition of the total sick depends in part on the size of the healthy group relative to the sick. The larger the healthy group, the more it will alter the sick—toward frailty or toward robustness. If the healthy group is very large, then proportionally small decrements from the healthy group may alter its composition little while changing sick group composition substantially.

Second, it seems plausible that, while the healthy group should mimic other binary-frailty closed groups in having at most one absolute deceleration, the sick group may decelerate more

than once. This possibility arises from the interaction of two flows changing sick group composition—the decrement to mortality and the increment from the healthy—and the fact that the direction of the effect of the latter increment may change as the composition of both groups changes.

To explore these possibilities, I simulate the model under a wide range of parameter values.

Simulation Procedure

The purpose of the simulations is to compare the age at deceleration among sick and healthy groups who age according to the model just described. Each cohort is constructed as a set of two multistate life tables—one representing the frail subgroup, and one the robust—which never exchange members. The age-specific mortality of the sick and healthy groups is then constructed as the average mortality of their respective frail and robust members, weighted by the size of each subgroup. In constructing the life tables, I use standard formulas which assume constant mortality within each age interval (Preston et al. 2001: 46-47). Since this assumption is inconsistent with the model assumption of Gompertz mortality within each subgroup, I use age intervals of only one-tenth of one year, ranging over 100 yearly ages (producing 1001 age observations). All frail members in all cohorts are extinct by the end of the age interval, ensuring that all decelerations are observed.

The onset of mortality deceleration is defined as the point when the second derivative of mortality becomes negative. The first derivative of mortality is estimated as the two-sided average difference in mortality, and likewise, the second derivative is estimated as the two-sided

average difference in the first difference.[7] Since this measure requires comparison points on each side of each age, the second derivative cannot be estimated for the first two and last two age units, leaving 997 age observations with estimated second derivatives of mortality for each group.

Since the model may be sensitive to the values of many parameters, it is desirable to evaluate specifically those groups whose mortality seems realistic for human populations. Yet many of the parameters of interest are putatively unobserved in empirical data, and so it is difficult to directly evaluate whether they are reasonable or realistic. The strategy I adopt, then, is to consider primarily those sets of groups that generate *aggregate* cohort mortality that resembles aggregate mortality of real cohorts in the Human Mortality Database (HMD).[8] To do this, I first allow the parameters to covary freely over a fairly wide range of values, producing 2,603,664 total simulated cohorts, each containing groups defined by frailty and robustness, sickness and health.[9] I then estimate a Gompertz model on the full cohort produced by each set of groups, and

---

[7] In other words, the slope of mortality at age *a* is estimated as half of the difference between the mortality at age *a+1* and the mortality at age *a-1*. I use the two-sided difference because I consider it to be the best measure for equivalent purposes in empirical research, since one-sided differences may exaggerate distortions arising from tempo effects (by simultaneously inflating and depressing mortality at adjacent ages, when more or fewer people than expected die during a given interval), while two-sided differences smooth over such distortions. However, since the model here is deterministic (in the sense that there is no stochastic variation that might produce such tempo distortions), the choice between two-sided and one-sided difference estimates of the derivatives is arbitrary in this context.

[8] 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). I use all cohort (vs. period) data included in the HMD.

[9] Specifically, I run the models on all combinations of these parameter ranges: α_μ=[.0031, .0101] in units of .001, as well as .00001, .00005, .0001, .0005, .0101, and .0201; β_μ=[.05, .125] in units of .015; f_μ=2, 3; m=2, 3; α_ω={.00001, .00005, .0001, .0005, .0101, .0201}; β_ω=[.05, .125] in units of .015; f_ω=2, 3; $\pi(0)$ for both groups =[.3,.9] in units of .1; and the baseline proportion healthy =[.3,.9] in units of .1. These parameter values by selecting an initial set that previous research with a closed-group model (Wrigley-Field 2013) suggested might yield realistic aggregate values, and then were iteratively augmented (i.e., extended through trial-

an identical model on each cohort in the HMD.[10] Finally, I construct a parallelogram around the estimated Gompertz intercept and slope of the real data (a shape that, consistent with Strehler and Mildvan's [1960; see also Finkelstein 2012; Zheng, Yang, and Land 2011] observation, fits well the distribution of these data), and keep only the 46,684 simulated cohorts whose aggregate slopes and intercepts fall within the parallelogram of the HMD slopes and intercepts. Of those simulated cohorts, 11,639 have a deceleration among both the healthy and the sick, allowing comparison of the age at each.[11] These 11,639 comprise the main sample in the paper, because, following the standard practice in empirical research on deceleration, the main object of interest is the direction and magnitude of the difference between the age at sick and at healthy deceleration. Since this measure fails to draw information from cohorts in which only one, or neither, group decelerates, this outcome should not be considered a comprehensive look at deceleration dynamics in an open-group model, but rather an example of what empirical

---

and-error) to find more cohorts that fell close to the HMD data. No parameter choices were made with regard to their likely deceleration outcomes.

The range of values of the baseline proportion frail at each group is chosen to exclude only groups whose baseline frailty is too low to make any deceleration likely. Previous research (Wrigley-Field 2013) shows that, in the closed-group version of this model with baseline frailty set to 75 percent, absolute decelerations tend to occur when the frail are between about one-third and two-thirds of the cohort.

[10] For purposes of this comparison, I consider each model life table to range over the ages 50 to 150. This choice does not affect the life table calculation, since the life tables are calculated with the baseline age rescaled to zero, but it does affect the interpretation of the parameters. The model therefore assumes that each disaggregated subgroup has Gompertz mortality from mid-life onward, but makes no assumption about earlier life. Likewise, the model assumes that a relatively high percentage of each group is frail at age 50. This high percent frail at age 50 does not impose an assumption of unduly high frailty at birth because, on the model given here with lifelong Gompertz mortality at the individual level, mortality selection overwhelmingly occurs at higher-mortality older ages. (The reasonableness of this assumption is discussed in Wrigley-Field [2013].) When I estimate the aggregated Gompertz models in the model cohorts and the HMD cohorts, I restrict the age range to 50-100 to match the availability of real data.

[11] Of the 46,484 cohorts, 26,570 (57%) show no deceleration in either group; 4,887 (10%) show deceleration only among the sick; 3,588 (8%) show deceleration only among the healthy; and so the 11,639 with decelerations among both groups—the cohorts analyzed in this paper—are 25% of the total.

researchers would conclude if they applied standard closed-group reasoning in an open-group context.[12] Figure 2 shows the aggregate Gompertz parameters of these simulated cohorts.[13]

Finally, to provide a basis for comparison, I construct the two alternative models. First, I construct a set of closed-group models based on the subset of "realistic" open-group models. For each set of mortality parameters included in that set of open-group models, I construct a closed-group model with those mortality parameters, but with the parameters describing the flow from sickness into health set to zero (in other words, the equivalent of the model for women and men given in Equations 1 and 2). This produces a set of 1,307 closed-group models, of which 181 show a deceleration among both groups.[14] These 181 closed-group cohorts are compared to the open-group cohorts in which the healthy and sick decelerate.

Second, I construct a set of models in which the frail and robust have equal rates of transitioning to sickness—that is, the frailty multiplier on sickness, f_{ω}, is set equal to 1—although the frail have higher mortality in each sickness group. I create one such model for each combination of parameters (except the frailty multiplier on sickness) in the main set of "realistic" models, producing 24,083 comparison models with equal sickness rates. Of these, 4,373 show a deceleration among both groups, and are compared to the main models.

---

[12] The outcome used here also loses some information to left truncation (if we do not think of baseline age as birth): some groups already have a negative second derivative of mortality at the baseline age. This is most common among sick groups: of the 11,639 cohorts considered here, 2,125 of the sick groups, and only 10 of the healthy groups, have already decelerated at the baseline age. Indeed, among cohorts in which the sick decelerate twice, the first deceleration overwhelmingly appears at baseline age. However, the order of deceleration is truncated only in the ten cohorts in which both groups have a negative second derivative at the baseline age.

[13] No simulations in the final set of simulated cohorts correspond to the upper left region of the HMD parallelogram, which corresponds to the most recent cohorts. This region is largely populated by simulated cohorts, but, at the parameter combinations used here, those cohorts do not show decelerations among both groups.

[14] The number of closed-group models—1,307—is substantially smaller than the number of HMD-compatible open-group models—46,684—because the latter include many sets of cohorts that share mortality parameters but vary in sickness parameters. Each such set of open-group models corresponds to a single closed-group model.

Simulation Results

In all models, the main outcome of interest is the sign (i.e., deceleration order) and magnitude of the difference between the sick and healthy groups' ages at deceleration. To provide a basis for comparing the open-group model results, which are the results of main interest, I begin by examining the 181 closed-group models in which both the healthy and the sick decelerate. As expected, in all cases the higher-mortality sick decelerate at a younger age than the lower-mortality healthy. As summarized in the first row of Table 1, the difference between the age at healthy deceleration and the age at sick deceleration ranges from 7.5 to 22.1 years, averaging 12.7 years. (Table 1 also reports median differences; none of the age difference distributions are badly skewed.)

The key quantity of interest in this paper is the difference between the healthy and the sick age at deceleration in the open-group model in which the frail are more likely to become sick, as well as to die. In short, as summarized in Table 1, in this model, it is no longer the case that the sick always decelerate at younger ages than the healthy; nor is the reverse reliably the case.

An immediate complication is that, of the 11,639 cohorts in the open-group model in which both groups decelerate, 1,425 cohorts (12%) show two distinct intervals of deceleration among the sick, with a return to accelerating mortality in between. To avoid any distortion in the results arising from comparing two deceleration points among the sick to one among the healthy, Table 1 reports separately the difference between healthy and sick deceleration among cohorts in which the sick decelerate only once; first sick decelerations compared to the single healthy

deceleration, among cohorts in which the sick decelerate twice; and, among those cohorts, second sick decelerations compared to single healthy decelerations.

Among those 10,214 simulated cohorts in which the sick decelerate a single time, as shown in the second line of Table 1, the difference between the age at healthy deceleration and the age at sick deceleration ranges from −36.8 to 22.6 (averaging −4.6). In other words, on average, the healthy decelerate at a younger age than the sick, with the sick decelerating first in only 27 percent of cohorts. This is the opposite of what we would expect from a naive extension of the closed-group generalization.

The third and fourth lines of Table 1 compare healthy and sick ages at deceleration when the sick decelerate twice. Among these cohorts, comparing the first sick deceleration to the single healthy deceleration, we find that the difference in ages ranges from −14.8 to 33.9 (averaging 11.4). Thus, here, too, the groups may decelerate in either order, although the first sick deceleration occurs before the single healthy deceleration in 99% of cohorts. Comparing the second sick deceleration to the single healthy deceleration, we find that the difference between healthy age and sick age at deceleration ranges from −39.3 to 3.8 (averaging −10.6, or median -8.4; this age difference has the most skewed distribution). The positive maximum difference shows that, in the open-group model, it is possible for the sick group to decelerate twice before the healthy group has decelerated. But 98% of second decelerations among the sick do occur after the healthy have decelerated. In short, when the sick decelerate twice, their decelerations nearly always fall on either side of the healthy group's deceleration—but it is possible, instead, for both sick decelerations to fall before or after the healthy deceleration.

An open-group model with equal sickness rates

For comparison to the main model used in this paper, consider the open-group model that drops the assumption that the frail become sick at a higher rate than the robust. The last three lines of Table 1 summarize the deceleration order among the 4,373 cohorts with equal sickness rates between the frail and robust in which both the healthy and sick groups decelerate. When the sick decelerate only once, 97 percent of the time, this deceleration occurs before the healthy deceleration, by 8.5 years on average. When the sick decelerate twice, in all cohorts considered here, the sick decelerate for the first time before the healthy, and for the second time afterward. These cohorts are rare: in the model with equal frail and robust transitions to sickness, only 14 cohorts (0.3%) have two sick decelerations.

In summary, it is not necessary for the frail to transition between groups at a higher rate than the robust to have mortality deceleration patterns in an open-group model that are different from those of the equivalent closed-group model in two respects: the healthy may decelerate at a younger age than the sick, and the sick may decelerate twice. But when the frail and the robust transition between groups at an equal rate, such patterns are very rare in the cohorts modeled here.

How predictable is open-group deceleration order?

The results presented so far demonstrate that the open-group model considered here does not conclusively predict the order of deceleration between groups, as the closed-group model does. We might also wonder to what extent the order of deceleration fluctuates with small changes in parameter values. If such fluctuations are rare, then deceleration order might still be

useful in an open-group context, as long as parameter values can be estimated (from data) or assumed (in the case of latent parameters) with reasonable precision. But if deceleration order fluctuates wildly with variation in parameter values, then it is unlikely to be useful as a test of an assumed selection model.

As a first step toward exploring the sensitivity of deceleration order to parameter values in the open-group model, I focus on the relationship between the baseline percent healthy and the order of deceleration between the healthy and the sick. I choose the baseline percent healthy as the parameter of interest because, unlike parameters that are specific to frail or robust subgroups, it is observable in empirical data—hence a potentially useful basis for empirical predictions of deceleration order in an open-group context. Moreover, as explained in the "Intuition and Hypotheses" section above, I expect the baseline percent healthy to play an important role in amplifying the selection dynamics associated with the movement from health to sickness.

To analyze the relationship between baseline percent healthy and deceleration order, I construct 7,961 sets of at least two cohorts that share all parameters except their baseline percent healthy and examine deceleration order within those sets. (When the sick group decelerates twice, I use the first sick deceleration.) In eleven percent of these sets, the order of deceleration switches, with the healthy decelerating first at some values of baseline percent healthy, and the sick decelerating first at other values.[15]

---

[15] Because this analysis is restricted to cohorts in the subset with similar aggregate parameters to the HMD cohorts, the sets contain varying numbers of cohorts, with a maximum of seven (corresponding to seven modeled values of the baseline percent frail, ranging from .3 to .9). The frequency of deceleration order switching given—eleven percent—corresponds to sets containing at least two cohorts. Because the sets are modeled over only a partial range of the proportion frail, and are further limited to values generating aggregate cohorts compatible with the HMD, I believe that the measure used here of the frequency of order switching is conservative.

Such deceleration order switching occurs at all measured values of the baseline percent healthy, but it tends to occur at very high values—86 percent on average. When the order of deceleration does switch within a set of cohorts sharing all parameters except the baseline percent healthy, in 94 percent of cases, the cohorts with the lower baseline percent healthy have the sick group decelerating at a younger age, while those with the higher baseline percent healthy have the healthy group decelerate first. To the extent that delayed mortality deceleration among the sick reflects frailty replenishment from the healthy, it is sensible that a larger healthy group would yield a later age of mortality deceleration among the sick.

In one set of cohorts, the order of deceleration switches twice: the healthy decelerate before the sick when the baseline proportion healthy is 0.8, and the sick decelerate first when it is larger or smaller.[16] To further explore the contribution of the baseline percent healthy for this set of cohorts, I model cohorts with this set of parameter values at baseline healthy values ranging fully—from half a percent to 99.5%—in units of half of one percent. Figure 3 shows the age at sick deceleration for these cohorts over the baseline percent healthy. The horizontal line at age 86.8 represents the age at healthy deceleration (constant across the cohorts, since mortality deceleration among the healthy is unaffected by the size of the healthy group). At values of the baseline percent healthy below 72.5% and above 90%, the sick decelerate before the healthy. At values of the baseline percent healthy between 72.5% and 90%, the healthy decelerate first. Such

[16] These cohorts have aggregate Gompertz parameter values similar to mid-19th Century Iceland and late-19th Century France. Their subgroup parameter values are: α_{μ}=.0031; β_{μ}=.08; f_{μ}=3; m=2; α_{ω}=.0001; β_{ω}=.11; f_{ω}=2; $\pi(0)$=.9.

sensitivity of deceleration order to the baseline percent healthy requires fairly precise consideration of the latter to interpret the former.[17]

Conclusion

There is a long tradition in demography of understanding that population composition alters and changes the interpretation of aggregate population behavior, dating at least back to Vaupel et al. (1979) and Vaupel and Yashin (1985). Most of this research has taken place in a closed-group context, and we now understand a great deal about the properties of population composition and selection in that context. There also is a smaller body of work on population dynamics that incorporates the insights of this classic body of mortality selection research (e.g., Manton et al. 1994, 1995; Woodbury and Manton 1983), and this area should expand in the future, with the explosion of longitudinal health data. But this study suggests some unforseen difficulties in integrating open-group models with one of the most important indicators of cohort heterogeneity: mortality deceleration.

The results in this paper demonstrate that the deceleration dynamics of open groups do not lend themselves to simple generalizations. Among closed groups, the higher-mortality group always decelerates at a younger age than the lower-mortality group, all else being equal. In open groups, however, either deceleration order is possible. Moreover, the order of deceleration can be sensitive to relatively small changes in parameter values, as in the perverse set of cohorts

[17] Since the sets here consider only cohorts with aggregate parameters similar to the HMD, order switching over values of the baseline percent healthy may appear less prevalent than it would with a more permissive universe of cohorts, and hence, larger sets of cohorts for comparison. Sets of cohorts can switch deceleration orders twice only if they contain at least three cohorts. 753 sets contain only two cohorts. Of these, 13 percent have a single order switch.

illustrated in Figure 3, in which the deceleration order switches back and forth as the baseline percent healthy increases.

The lack of reliable generalizations about the order of mortality deceleration in an open-group context poses a challenge to the usefulness of measurements of deceleration order. This, in turn, renders less promising one of a relatively small number of methods for testing selection theory and drawing inferences about unobserved heterogeneity in mortality risk.

Deceleration order is informative when it violates the predictions made by selection theory. Such a violation tells analysts that one of two things is going on. One possibility is that, although the deceleration arises from selection, the ceteris paribus conditions for the predicted deceleration order are not met, i.e. there is some other difference in heterogeneity between the groups (as Lynch et al. [2003] conclude about African-Americans and white Americans). The other possibility is that deceleration does not arise solely or primarily from mortality selection, but rather from deceleration in individuals' rate of aging, as prominent biodemographers have argued (e.g., Mueller et al. 2011; and see reviews in Vaupel 1997, Wachter and Finch 1997). Conversely, deceleration orders that fail to violate the predictions of selection theory are often considered to lend some support to selection theory (e.g., Horiuchi and Wilmoth 1998).

Such inferences are possible only when selection theory generates a clear prediction about which group should decelerate at an earlier age than another. The results in this paper demonstrate that, when we turn from comparing closed groups, such as birth cohorts, to open groups, such as health statuses, then—particularly when frailty may increase (or decrease) the risk of transitioning between groups—selection theory no longer makes such a clear prediction.

This renders the order of deceleration uninformative about heterogeneity without reliable information about underlying parameter values.

This paper provides a preliminary consideration of deceleration order between open groups. An important question that the paper does not address is whether these results generalize beyond the binary frailty model used here to continuous frailty models, particularly the widely-used gamma-Gompertz model (Gampe 2010, Horiuchi and Wilmoth 1998, Missov and Finkelstein 2011, Steinsaltz and Wachter 2006, Vaupel et al. 1979). Future research should also explore the sensitivity of these results to the functional form, over age, of the risk of becoming sick. Whether or not open-group deceleration dynamics turn out to be sensitive to model choice, however, these results show that, in the binary context, they are sensitive to latent parameter values. A practical implication is that, if deceleration order in an open-group setting is used either to test selection theory or to derive conclusions about cohort heterogeneity, the sensitvity of the conclusions to model assumptions should be explicitly modeled. It may be the case that, in specific open-group settings, such sensitivity analyses would reveal that some conclusions from deceleration order are well-supported. Based on the results presented here, such conclusions should not be assumed to be warranted in general.

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.

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.

Curtsinger, J.W., Fukui, H.H., Townsend, D.R. & Vaupel, J.W. 1992. Demography of genotypes: Failure of the limited life-span paradigm in Drosophilia-Melanogaster. Science, 258, 461-463.

Drapeau, M.D., Gass, E.K., Simison, M.D., Mueller, L.D., Rose, M.R., 2000. Testing the heterogeneity theory of late-life mortality plateaus by using cohorts of Drosophila melanogaster. Experimental Gerontology, 35 71-84.

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.

Finkelstein, M. 2012. Discussing the Strehler-Mildvan model of mortality. Demographic Research, 26(9): 191-206.

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

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

Horiuchi, S. & Coale, A.J. 1990. Age patterns of mortality for older women: An analysis using the age-specific rate of mortality change with age. Mathematical Population Studies, 2(4): 245-267.

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. 1992. Frailty and Survival. Genus, 47(3-4): 101-118.

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.

Manton, K. G., Stallard, E., Woodbury, M. A. & Dowd, J. E. 1994. Time-varying covariates in models of human mortality and aging: Multidimensional generalizations of the Gompertz. Journals of Gerontology, 49, B169-B190.

Manton, K. G., Woodbury, M. A. & Stallard, E. 1995. Sex differences in human mortality and aging at late ages: The effect of mortality selection and state dynamics. Gerontologist, 35, 597-608.

Missov, T. I. & Finkelstein, M. S. (2011). Admissible mixing distributions for a general class of mixture survival models with known asymptotics. Theoretical Population Biology, 80(1): 64-70.

Mohtashemi, M. & Levins, R. 2002. Qualitative analysis of the all-cause Black-White mortality crossover. Bulletin of Mathematical Biology, 64, 147-173.

Mueller, L.D., Rauser, C.L., & Rose, M.R. 2011. Does aging stop? Oxford: Oxford University Press.

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.

Rauser, C.L., Abdel-Aal, Y., Sheih, J.A., Suen, C.W., Mueller, L.D., & Rose, M.R. 2005. Lifelong heterogeneity in fecundity is insufficient to explain late-life fecundity plateaus in Drosophila melanogaster. Experimental Gerontology, 40 (8-9): 660-670.

Robine, J.-M., Jagger, C., Mathers, C.D., Crimmins, E.M., & Suzman, R.M. (eds.) 2003. Determining Health Expectancies. Hoboken: John Wiley & Sons.

Rogers, A. 1992. Heterogeneity and selection in multistate population analysis. Demography, 29, 31-38.

Steinsaltz, D. 2005. Re-evaluating a test of the heterogeneity explanation for mortality plateaus. Experimental Gerontology, 40(1-2): 101-113.

Steinsaltz, D.R., & Wachter, K.W. Understanding mortality rate deceleration and heterogeneity. Mathematical Population Studies, 13(1): 19-37.

Strehler, B.L., and Mildvan, A.S. (1960). General theory of mortality and aging. Science 132(3428): 14-21.

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. and J. R. Carey. 1993. Compositional interpretations of medfly mortality. Science, 260: 1666-1667.

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.

Wachter, K. W. & Finch, C. E. (eds.) 1997. Between Zeus and the Salmon: The Biodemography of Longevity. Washington, D.C.: National Academy Press.

Wrigley-Field, E. 2013. Mortality Deceleration and Mortality Selection: Three unexpected implications of a simple model. CDE Working Paper Series #2013-01. http://www.ssc.wisc.edu/cde/cdewp/2013-01.pdf

Woodbury, M. A. & Manton, K. G. 1983. A mathematical model of the physiological dynamics of aging and correlated mortality selection. 1: Theoretical development and critiques. Journals of Gerontology, 38, 398-405.

Zeng, Y., Crimmins, E.M., Carrière, Y., & Robine, J.-M. (eds.) 2006. Longer Life and Healthy Aging. New York: Springer.

Zheng, H., Yang, Y., & Land, K.C. 2011. "Heterogeneity in the Strehler-Mildvan General Theory of Mortality and Aging." Demography, 48: 267-90.

TABLES

Table 1. Age at deceleration in healthy vs. sick groups

Model	Sick decel.	Difference in deceleration age in healthy vs. sick groups (in years)				Proportion of cohorts in which sick group decelerates before healthy group
		Min	Max	Mean	Median	
Closed-group model (n=181)	Single (n=181)	7.5	22.1	12.7	11.6	1
Main model: Open-group model *with* frailty multiplier on sickness >1 (n=11,639)	Single (n=10,214)	-36.8	22.6	-4.6	-4.4	.27
	First (n=1,425)	-14.8	33.9	11.4	10.6	.99
	Second (n=1,425)	-39.3	3.8	-10.6	-8.4	.02
Open-group model *without* frailty multiplier on sickness >1 (n=4,373)	Single (n=4,357)	-4.5	52.8	10.9	9.5	.97
	First (n=14)	8.5	12.4	11.4	11.7	1
	Second (n=14)	-6.4	-3.1	-5.1	-5.1	0

Note: All results in Table 1 are for simulated cohorts in which both healthy and sick groups decelerate. In all cases, the healthy group decelerates only a single time, so all sick decelerations (single, first, or second) are compared to the single healthy deceleration in the cohort.

FIGURES

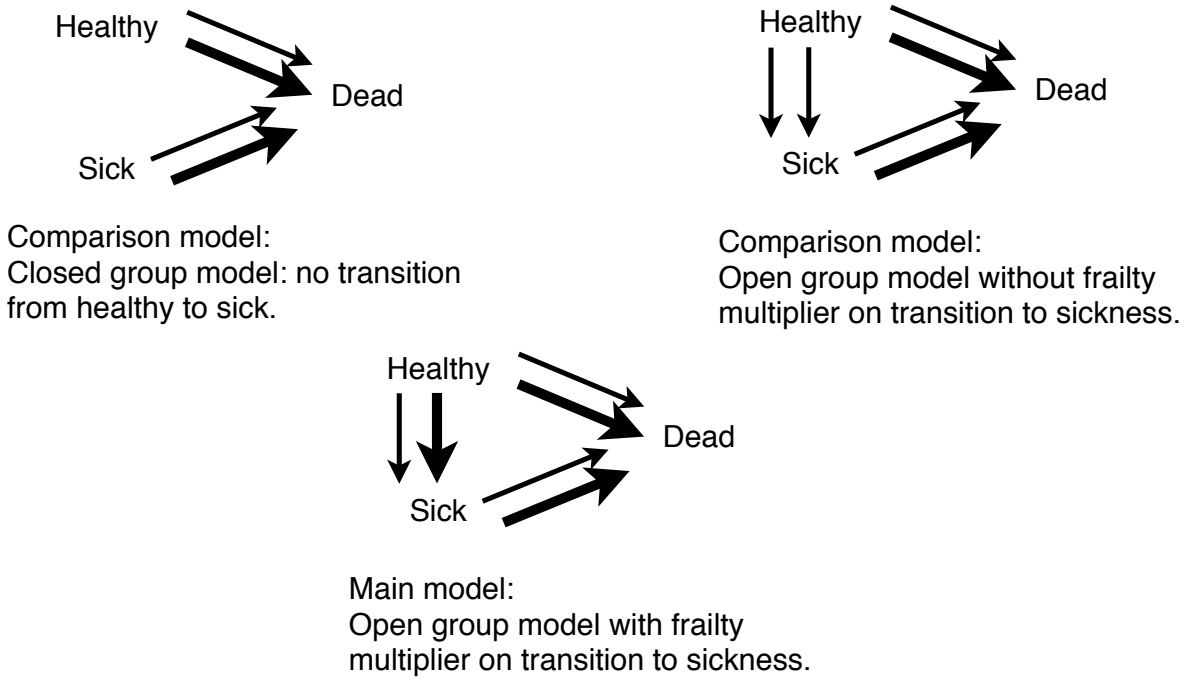

Figure 1. Graphical representation of the three models used in the paper.

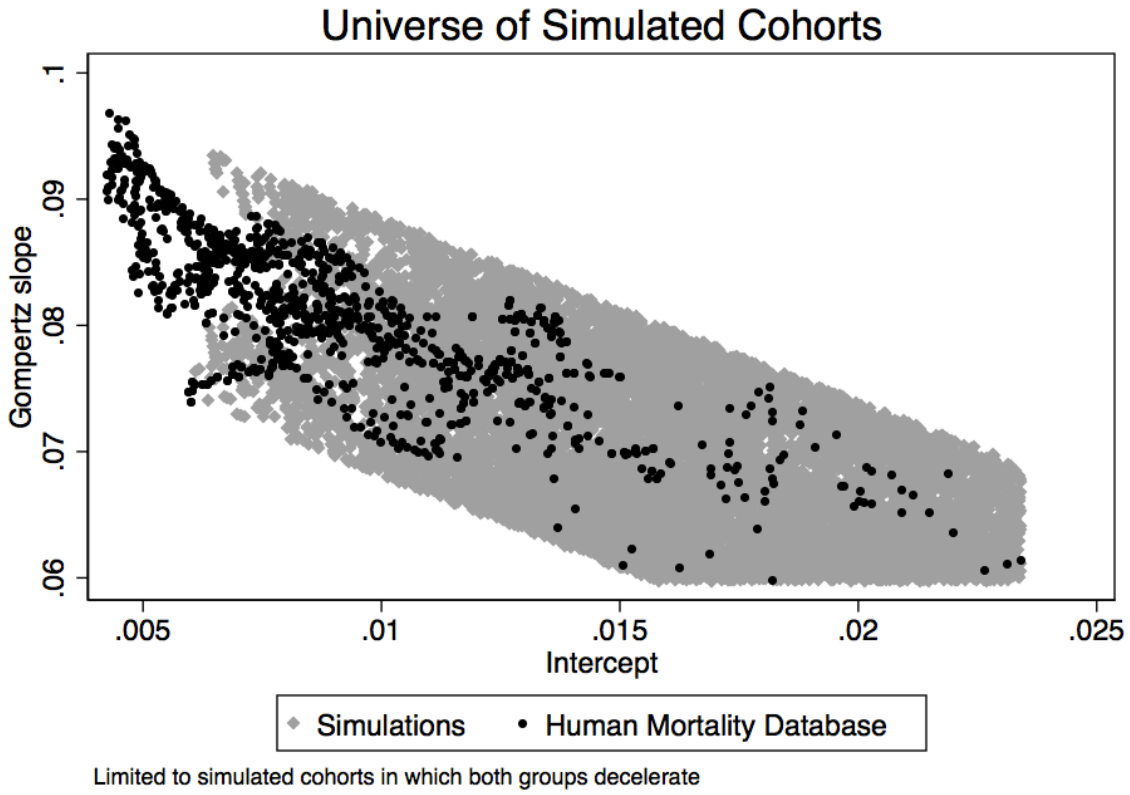

Figure 2. Aggregate parameters of the universe of simulated cohorts and the Human Mortality Database cohorts.

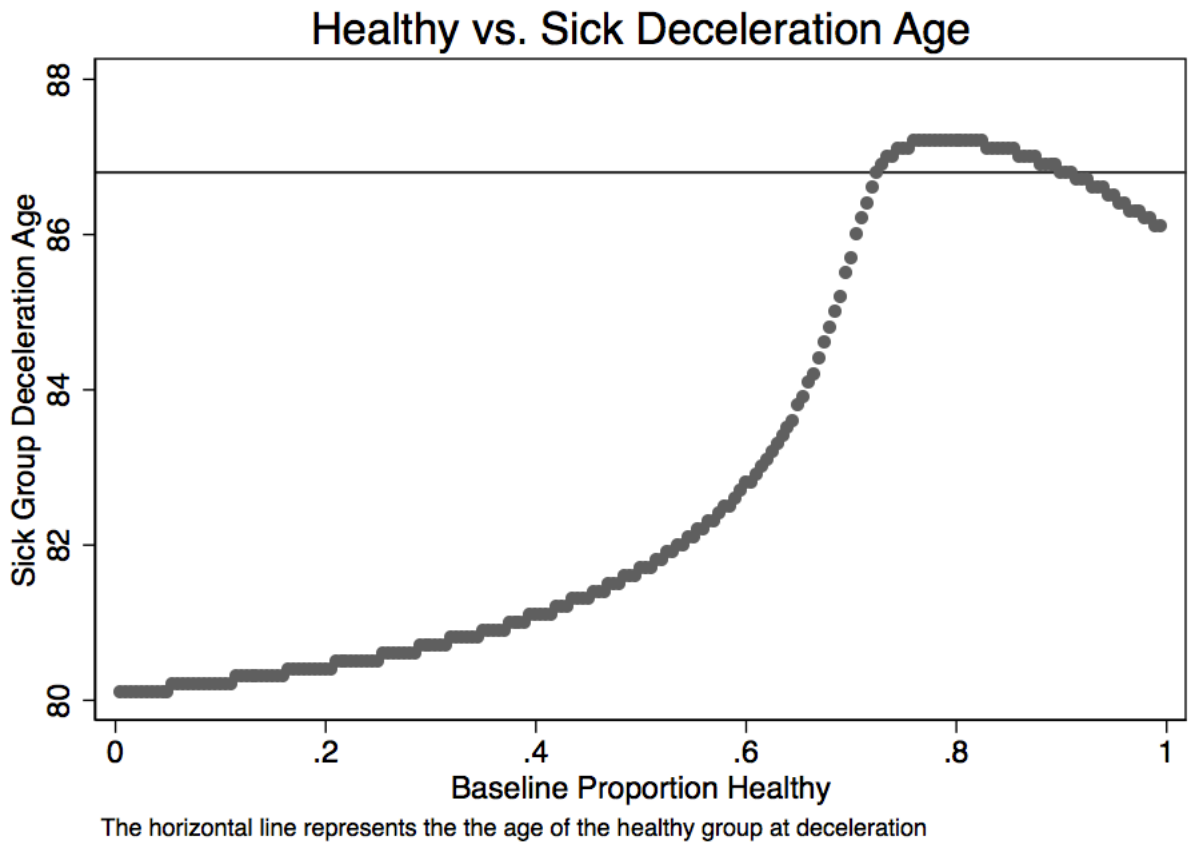

Figure 3. Age at mortality deceleration among the sick group as a function of the baseline percent healthy, compared with age at deceleration among the healthy group.

CHAPTER 3:

Mortality Selection or Life Course Changes?

The Missing Explanation for the Black-White Mortality Crossover in the United States

Introduction

The black-white mortality crossover is the well-known phenomenon that aggregate black mortality is higher than white mortality from birth into old age but drops below white mortality at the oldest ages (for recent cohorts, around age 85). It has become commonplace in recent demographic writing to assert that there are but two explanations for the crossover: age misreporting and mortality selection (e.g., assertions in Dupre et al. 2006, Johnson 2000, Lynch et al. 2003, Masters 2013, Nam 1995, Parnell and Owens 1999, Sautter et al. 2012, Yi and Vaupel 2003). The age misreporting hypothesis is that the crossover is an artifact of elderly black ages being overstated in some datasets (e.g., Coale and Kisker 1986; Elo and Preston 1994; Preston et al. 1996). The mortality selection hypothesis is that the crossover reflects that, given the conditions of stark deprivation for many African-Americans, only the hardiest, most robust blacks survive to old age (e.g., Vaupel et al. 1979, Vaupel and Yashin 1985). As data quality has improved and the age misreporting hypothesis is increasingly discounted, the selection hypothesis is often accepted by default.

Yet the original tradition of research on the crossover sometimes engaged a third explanation: that biological and social factors led black disadvantage in mortality to diminish over the life course at the individual level, not only in the aggregate (e.g., Berkman et al. 1989,

Corti et al. 1999, Manton et al. 1979, Wilmoth and Dennis 2007). Our contention is that demographic research must not forget the life course hypothesis.

In particular, we show that failing to consider the life-course alternative to the selection hypothesis has led some accounts of the crossover, which really are life course accounts, to be misdescribed as instances of mortality selection. This, in turn, has distorted demographic understanding of how mortality selection works. Therefore, in this chapter, we clarify the conditions under which the crossover would result from mortality selection and the conditions under which the crossover would result from life course changes. What is at stake in the distinction between selection and life course explanations of the mortality crossover is accurately describing the persistence and life course patterning of racial inequality in the United States.

The stakes in distinguishing mortality selection from diminishing black disadvantage

The crossover is only an extreme case of a fundamental problem afflicting the status of mortality as evidence about either social inequalities or biological processes of aging. The problem is that mortality patterns at older ages may reflect both changes experienced by individuals and changes in group composition, wrought by mortality selection. When the two are conflated in the mortality evidence, but only one is explicitly considered, conclusions may be mistaken. A classic, but surprisingly undiscussed, example shows the stakes in differentiating mortality selection from declining disadvantage by suggesting that the black-white mortality crossover may have distorted demographic understanding of slavery.

Robert William Fogel's work on nutritional status and mortality famously reshaped demographic debates about the origins of historical mortality declines. Fogel (1986) compares

the mortality of black slaves and free whites in the antebellum U.S. South, and shows that black excess mortality is very large in early childhood and very small afterwards, disappearing entirely by age 20. Fogel concludes that "it was excess death rates of slave children under 5 that accounted for the difference between the overall death rates of U.S. slaves and U.S. whites during the late antebellum era. Moreover, the fact that U.S. slaves and whites had similar life expectancies after age 20 suggests that it was not the general virulence of the disease environment but conditions specific to young children." (Fogel 1986: 51)

Fogel's readers, however, might wonder might wonder whether Fogel's data might not equally or perhaps better support an alternative conclusion: that slavery is so consequential for survival that slaves, to survive to age five, had to be enormously hardy—so much so that even the ongoing deprivations of slavery did not raise their mortality above the rate for whites. Modern readers, that is, might wonder whether Fogel's result is an artifact of mortality selection. But what are we to conclude when our priors are not so powerful that they force one interpretation or the other?

The crossover therefore crystallizes the challenge that mortality selection poses to studies of mortality more generally: inequalities in contemporaneous conditions (e.g., the conditions of life among elderly blacks and whites today, or among black slaves and free whites in the antebellum South) are conflated with the legacy of inequalities in the past, each pushing the difference in the mortality hazard in different directions. As a measure of current inequality, mortality differentials are inherently biased. To use mortality crossovers as evidence either of

diminishing inequalities in a cohort's older ages, or alternatively, of consistent inequality at its younger ages (mortality selection), requires a principled basis for differentiating the two.[1]

Today, research—at least research framed as studying the mortality crossover—is more likely to make the opposite assumption. Where Fogel saw only diminishing black disadvantage, more recent demographic work sees only mortality selection. We argue that this, too, may be an error for understanding the true patterning of black disadvantage in survival over the entire life course. Our argument begins by reviewing the three potential explanations for the crossover.

The Age Misreporting Explanation of the Crossover

From the start, demographers voiced concerns that mortality crossovers may result from measurement error (Pearl 1922). Specifically, various early inquiries established that African American ages at death are systematically exaggerated at the oldest ages on death certificates, and that age exaggeration was more prevalent among African Americans than among whites (Hambright 1968; Rosenwaike 1981). Consequently, official age-specific death rates, which draw on death-certificate data, may inappropriately compare relatively younger African Americans to relative older whites, thus creating negative bias in black-white mortality differentials. Subsequently, various authors argued that the black-white mortality crossover in official vital statistics is entirely explained by age misreporting (Coale and Kisker 1986; Elo and

[1] The general problem of distinguishing mortality selection from individual patterns also occurs far more widely than the black-white mortality crossover. As a general pattern, mortality differentials tend to diminish at older ages (Wilmoth and Dennis 2007: 298). A recent example of the difficult epistemic situation researchers are placed with, without a principled basis for distinguishing individual-level patterns from selection, concerns Holocaust survivors. Showing that some of those who survived the Holocaust have lower mortality than those whose birth timing let them narrowly avoid it, Sagi-Schwartz et al. (2013) hypothesize that surviving a traumatic experience yields health-protective coping skills—though they note that mortality selection might also explain their results. Our contention in this article is that leaving it to the whim of researchers whether they find such individual-level hypotheses persuasive or not is a dissatisfying foundation for population science.

Preston 1994; Preston et al. 1996). Some of this work, however, relied on incomplete corrections. For example, Preston et al. (1996) compared age-corrected African American death rates to uncorrected white death rates. When Hill et al. (2000) corrected for both black and white age misreporting at old ages, they found that age misreporting does not entirely explain the black-white mortality crossover in official statistics: while the age at crossover moves upward after corrections, a crossover remains above age 90.

More recent demographic research has avoided the limitations of death certificate data by analyzing better data in which age is measured early during decedents' lifetimes. In these data, black-white mortality crossovers are unambiguously present. For example, Kestenbaum (1992) analyzed age-verified Medicare enrollment data and found a black-white mortality crossover in old age. Similarly, Dupre et al. (2006) and Sautter et al. (2012) analyzed ESPE data from North Carolina, which come from a longitudinal study where age is verified during respondents' lifetimes, and found unambiguous black-white mortality crossovers. In the wake of these findings, the present-day consensus appears to be that data errors do not fully account for the black-white mortality crossover.

The Mortality Selection Explanation of the Crossover

The mortality selection explanation is that the black-white mortality crossover results from the disproportionate selection by mortality of frailer members out of the black population, so that age-specific mortality comparisons compare a highly selected, robust group of surviving blacks to a less selected, less robust group of surviving whites. Since robust individuals have lower mortality than frail individuals, the compositional shift of the black group toward more robust

members compared to whites can push aggregate black mortality below white aggregate mortality in old ages. The key point of the pure selection model is that aggregate black and white mortalities can cross purely as a result of compositional shifts among the survivors as the population ages, even if black and white mortalities never cross at the individual level.

The canonical pure selection model with binary frailty posits the existence of four internally homogenous groups, where the mortality of every individual *I* is fully determined by her age, a_i, race (black vs. white) and frailty (frail vs. robust), as in Equation 1:

$$\begin{aligned} \mu_{w,r}(a_i) &= \alpha e^{\beta a_i} \\ \mu_{w,f}(a_i) &= f\alpha e^{\beta a_i} \\ \mu_{b,r}(a_i) &= b\alpha e^{\beta a_i} \\ \mu_{b,f}(a_i) &= bf\alpha e^{\beta a_i} \end{aligned} \tag{1}$$

In this model, individual-level mortalities $\mu(a_i)$ are Gompertz. Individual mortalities never cross (are log-parallel) because they share the same slope parameter, β, and differ only by a group-specific intercept. The shared component of the intercept is α, the intercept for robust whites. This intercept is multiplied by the fixed scalar $f > 1$ for frail individuals and the fixed scalar $b > 1$ for black individuals.

The aggregate mortality, $\bar{\mu}_k(a)$, of race $k=w,b$ at age *a* is the average mortality of all the black and white survivors at that age, i.e., the weighted average of the mortality of the robust and the frail, weighted by the proportion of the total survivors who are in each group, as shown in Equation 2:

$$\bar{\mu}_k(a) = \pi_k(a)\mu_{k,f}(a) + \big(1 - \pi(a)\big)\mu_{k,r}(a) \quad (2)$$

where $\pi_k(a)$ represents the proportion of the race that is frail at age *a*.

If group membership (with respect to race and frailty) is fixed, then the only process that causes the proportion frail to change over age is that the frail die faster than the robust. Moreover, because of higher black mortality at all ages, at any age, the cumulative loss of the frail will be larger among blacks than among whites. Thus, being black is simultaneously a circumstance that raises mortality in individuals, and, increasingly with rising age, an indicator of being robust, which lowers mortality in individuals. It is this disproportionate loss of the frail among blacks that can produce a mortality crossover (Manton and Stallard 1981, Vaupel et al. 1979, Vaupel and Yashin 1985).

This model represents the classic mortality selection explanation of the crossover. This article elaborates on when this explanation will hold, and when a crossover might instead be produced by life course changes in individuals. Before elaborating the life course hypothesis, we turn our attention to what exactly this *frailty* might be.

What is frailty?

How should we characterize the people "at highest risk of dying"—the *frail*, in the language of mortality selection theory? Early discussions of mortality selection (including discussions by some of its founding theorists) sometimes assumed that frailty and robustness signified genetic differences (e.g., Olshansky 1995, Manton et al. 1979). Yet what is needed for

some type of heterogeneity to play the role ascribed to frailty is not always met by genetic heterogeneity, and certainly is not only met by it. The essential precondition for mortality selection is that individuals are stably ordered in their risk of dying as they age: those that are more likely to die than others in, say, their 40s, must also be more likely to die in their 70s, assuming they live so long. Without some reasonable degree of stable ordering, there can be no mortality selection of the sort that could produce a crossover: the people at greatest risk of dying while young would disproportionately die young, but their absence would not lower the aggregate hazard of the survivors at older ages, since they would not have been high-mortality people at those ages had they lived. Anything that contributes to this stable ordering can be considered *frailty* in the sense that it plays the part demanded by mortality selection theory: frailty is distinguished by its structure, not its substance.[2]

One way to understand what this stable ordering condition means is to ask: in what circumstances would stable ordering not hold, and thus, mortality selection not occur? Here are three. First, no mortality selection would occur within black and white subpopulations if the racial subpopulations were homogeneous in their mortality risk. Second, no mortality selection

---

[2] In this connection, there is historical interest in this passage from Manton et al. (1979: 293): "There is, however, one further assumption which the selection model makes which is not only untenable, but which suggests another mechanism that could account for the crossover either by itself or in conjunction with selection. This last and perhaps most basic assumption is that selective pressures act on fixed genetic characteristics for longevity which remain constant throughout the life span. In effect, this suggests that the organisms' vitality at any given age is precisely determined by genetic endowment. […] It is the case, however, that the organism's physical well being is modified in interaction with the environment."

In this remarkable passage, coming at the very beginning of demographic discussion of mortality selection, Manton et al. identify something close to the structural requirement of the selection model (what we call the 'stable ordering condition'); they recognize that genes alone cannot meet this requirement because of the role that socioeconomic factors play in structuring mortality risk; but they take that to be a refutation of the selection model, never considering that those very socioeconomic factors might play the necessary structural role. (The Manton et al. [1979] paper remains one of the most creative and specific in the crossover literature, some 35 years after its publication.)

Passages like these underscore what an important shift it was in the mortality selection literature to begin to explicitly incorporate socioeconomic dimensions of heterogeneity into mortality selection research.

would occur if aging involved a kind of 'resetting the clock' on mortality risk, such that the people at higher risk of dying at a given age were not also at a higher risk of dying above that age. Formally, this might obtain if heterogeneity within each race group was an iid random draw at each age. Substantively, such a resetting might occur in several ways: for example, if early health insults caused one to adopt a healthier lifestyle; alternatively, it might occur if mortality risk were driven by mechanisms like mitochondrial DNA mutations, whose deleterious effects can strike suddenly at any age; or if the causes of death at young ages and older ages were sufficiently different that they did not share common risks (on this point, see especially Li et al. 2013, Horiuchi and Wilmoth 1998, Fenelon 2013, Manton et al. 1979). Third, no mortality selection would occur in a 'zombie counterfactual' in which mortality did not remove anyone from the population, allowing them to continue to experience mortality.

Although the second possibility may be partially substantiated, all three possibilities are clearly stylized (even if health risks do shift over age, there is a great deal of year-to-year continuity: it is certainly not the case that the mortality risks are memoryless from age to age). Therefore, the most relevant question may be not *whether* mortality selection occurs at all—it assuredly does—but rather, *how much* it matters, and, when it comes to explaining the crossover, whether it is the *only* thing that matters. This poses the question: what else might contribute to the crossover?

The Missing Explanation: The Life Course Explanation of the Crossover

In contrast to the selection explanation, the life course explanation stipulates that, even if the cohort retained its original membership, black mortality would fall below white mortality at

advanced ages. Under this hypothesis, black mortality falls below white mortality, not because mortality selection changes the composition of the black group more than the white group *in the aggregate*, but because something in the biological or social processes that constitute aging reduces—indeed, reverses—black disadvantage in survival *at the individual level*.

A pure life course model differs from the pure selection model of Equation 1 in two respects. First, it eliminates all heterogeneity within the black and white groups to eliminate selection. Second, it stipulates that black mortality start out higher than white mortality, but grow more slowly in individuals over age. Equation 3 gives one such pure life course model:

$$\begin{aligned} \mu_w(a_i) &= \alpha e^{\beta a_i} \\ \mu_b(a_i) &= b\alpha e^{c\beta a_i} \end{aligned} \tag{3}$$

with $b > 1$ and $0 < c < 1$. As before, a_i is the individual's age, α is the shared component of the intercept, and β is the shared component of the slope. In contrast to the pure selection model of equation 1, however, this model does not include any heterogeneity conditional on race (there is no frailty multiplier) and hence does not allow for changes in the compositions of the race groups as the population ages. Most importantly, the model stipulates that the black mortality starts out higher than white mortality, $b > 1$, and that black mortality increase more slowly than white mortality with age, $0 < c < 1$. In sharp contrast to the pure selection model, individual-level mortalities in the life-course model are not log-parallel (because blacks and whites have different slopes). Therefore, this model stipulates morality convergence in general, and, with suitably chosen parameter values, a black-white mortality crossover at some advanced age.

The life course explanation in the demographic literature

The defining feature of the life course model is that it permits black disadvantage relative whites change with age at the individual level. Several papers in the crossover literature explicitly acknowledge the life course possibility (e.g., Berkman et al. 1989, Corti et al. 1999, Manton et al. 1979, Wilmoth and Dennis 2007). One of the earliest to do so is a classic in the crossover literature. In discussing a black-white crossover among residents of public housing, Berkman et al. (1989) consider that blacks in public housing may be a deeply disadvantaged population heavily selected by mortality (mortality selection), that whites end up in public housing only through a series of life events predictive of mortality (selection by means other than mortality), and that the social networks in public housing may confer advantages for blacks and disadvantages for whites (a life course explanation).[3] Their discussion is an example of the simultaneous attention to selection and life course dynamics that has characterized the mortality selection literature at its best.

More generally, the demographic literature has identified several potential sociologically plausible mechanisms for such a reduction in black disadvantage over the life course. As a general matter, if mortality risk were decomposed into particular traits, then one (or both) of two kinds of changes might underlie the life course model: changes in the racial distribution of a trait (that raises or lower mortality) due to changes in individual status (not changes in population

[3] "Following this line of reasoning, we speculate that elderly blacks in public housing may experience a tightly knit community surrounded by family, friends, and religious groups that is protective of their well-being. Whites, on the other hand, may be much more isolated and may experience life in public housing as less supportive and more threatening. We consider the provision of a succinct explanation of these survivorship differentials in terms of social and economic conditions and/or possibly unmeasured […] physical and/or mental conditions as a major challenge for the future." (Berkman et al. 1989:676-77)

composition due to mortality selection), or changes in the effect of a trait (that is unequally distributed by race) on individuals' mortality.

An instance of the first kind of change would be a trait that always raises mortality, but becomes more prevalent for whites than blacks. For example, elderly blacks might experience a survival benefit compared to elderly whites because elderly blacks are less likely to live alone (Goldsheider and Bures 2003; Ruggles 1994), in the context that the elderly who live alone face substantially greater mortality (e.g., Luo et al. 2012, Steptoe et al. 2013).[4] Conversely, another instance of this kind of change would be a trait that always lowers mortality, but ceases to be disproportionately experienced by whites. A notable example is the proposal that black disadvantage is lower at old ages because everyone has access to Medicare (e.g., Sautter et al. 2012, Wilmoth and Dennis 2007). (Presumably this could not by itself cause black mortality to fall below white mortality, but could reduce the extent of black disadvantage and so contribute to a crossover alongside other changes.)

An instance of the second kind of change would be a trait that is prevalent among blacks, raises mortality at younger ages, and lowers mortality at older ages. For example, a body shape that is classified as overweight or obese is more common among blacks, and some studies have concluded that this status is harmful to health at younger ages but protective among elderly (the so-called "obesity paradox"; see, e.g., Curtis et al. 2005).[5]

---

[4] The authors cited in this sentence, as well as the Elwert and Christakis (2006) and Curtis et al. (2005) studies cited below, have documented the relevant facts about the racial distribution of the relevant traits, and/or their relationships to mortality; the inference that these relationships could contribute to a crossover at the individual level is our own.

[5] Of course, the obesity paradox is itself a mortality crossover, and as such it too might either reflect either a real change at the individual level or be an artifact of selection.

Either type of change might further contribute to a crossover if it also has different effects on mortality for each race. For example, Elwert and Christakis (2006) find that mortality among whites, but not blacks, rises sharply with widowhood, which is a relatively common experience among the elderly but rare among younger people in the United States. In another example, discussed further below, Dupre et al. (2006) argue that religiosity may be more protective for African-Americans than white Americans. We note that such traits whose effects might vary by race might effect a crossover in part because the causes of death (e.g., heart disease) that respond to the trait rise in importance with age, while other causes of death that disadvantage blacks (e.g., accidents and homicides) are concentrated at younger ages.

Finally, such potential sources of survival advantage might intersect with a biologically-based convergence, if the nature of biological aging is such that advantages tend to narrow over the oldest ages simply because death can be postponed for only so long (Wilmoth and Dennis 2007: 303).[6] Some analysts have proposed that the biological rate of aging may increase more slowly for blacks than for whites (Liu and Witten 1995), albeit with little motivation in the form of specific mechanisms of aging. In short, a wide variety of life course processes could in principle result in elderly white mortality overtaking elderly black mortality, not just via selection, but at the level of individuals.

As this discussion shows, the distinction between the mortality selection and life course explanations of the crossover is not a distinction between biological and social explanations.

---

[6] Although this argument is largely ignored in the context of the black-white crossover (apart from Liu and Witten (1995), reviewed below), many consider an argument like this to explain the old-age convergence in the hazards by socioeconomic status in the United States (e.g., Dupre 2007). Still further afield, demographic work on mortality deceleration, as opposed to mortality crossovers, has grappled with the possibility that such deceleration might result from selective mortality, from decelerating mortality at the individual level, or both (e.g., Horiuchi and Wilmoth 1998).

Indeed, either kind of explanation can be based in any combination of biological and social factors. Rather, the distinction concerns the level of aggregation at which some change produces the crossover.[7]

But many crossover articles, especially most recent articles (e.g., Dupre et al. 2006, Johnson 2000, Lynch et al. 2003, Masters 2013, Nam 1995, Parnell and Owens 1999, Sautter et al. 2012, Yi and Vaupel 2003), do not acknowledge the possibility of mortality crossover or convergence at the individual level. As we will show, this life course explanation has been written out of the crossover literature to the latter's detriment. The goal of this article is not to advocate that the life course explanation be accepted as true, but to advocate its inclusion in the established set of crossover explanations.

Consequences of Ignoring the Life Course Explanation

As an explanation of the black-white mortality crossover, mortality selection is often accepted by default once bad data has been ruled out (e.g., Dupre et al. 2006, Johnson 2000, Masters 2013, Nam 1995, Parnell and Owens 1999, Sautter et al. 2012, Yi and Vaupel 2003). We advocate that it be explicitly considered alongside the life course alternative in many substantive contexts.

There are three reasons to test the selection explanation, rather than assume it as true. First, we have shown that demographic work has identified some potential mechanisms for an

[7] An explanation could be both a life course explanation and a selection explanation if it operates at both levels of aggregation. For example, imagine that the black and white cohorts each contain two subpopulations: a higher-mortality subpopulation ("the frail"), which dies primarily of illnesses worsened by stress, malnutrition, and environmental toxins, and in which blacks have higher mortality than whites at all ages; and a lower-mortality subpopulation ("the robust"), which dies primarily of loneliness and social isolation, and in which blacks have higher mortality in young adulthood but, because they are less likely to live alone, have lower mortality at very old ages. Then mortality selection results in the survivors of both races being weighted toward the group in which there is a mortality crossover at the individual level.

individual-level crossover or convergence in the black and white hazards, implying that selection ought not be the *a priori* accepted explanation.

Second, there exists a diverse array of crossovers (not only the black-white crossover), and mortality selection may account for some but not others. For example, Thornton (2004) finds a mortality crossover between Navajo Native Americans and the total U.S. population; Hoffmann (2008) finds a mortality crossover among smokers and non-smokers in the United States; Guillot (2007) finds a mortality crossover among Russian and Kyrgyz infants in Kyrgyzstan; Huang and Wu (2010) find a mortality crossover between individuals with public health insurance and those without in China; and Yi and Vaupel (2003) find national mortality crossovers between Sweden, Japan, and the Han Chinese of China. Demography therefore needs strategies to empirically adjudicate the causes of crossovers in general. In the absence of such tests, whether a crossover is interpreted as reflecting selective mortality, or alternatively as reflecting life course processes in individuals, is determined not by any empirical evidence, but merely by the analysts' priors. This is an untenable basis for adjudicating hypotheses with very different implications about the extent and nature of inequality in mortality.

But perhaps most fundamentally, we advocate explicitly considering mortality selection explanations against alternatives because the habit of simply assuming that it accounts for the crossover has led the selection explanation itself to be badly distorted and misunderstood.

A striking example is Liu and Witten (1995), and the diametrically opposed ways that article has been summarized in the subsequent literature—even by its own authors. Liu and Witten's model is a life course explanation that has been repeatedly misunderstood as a selection explanation.

Liu and Witten (1995) offered an analysis of a crossover between two subpopulations, driven by an assumption that each subpopulation has a predetermined maximum lifespan. Specifically, the model, repeated in Liu et al. (2008), involves making two assumptions. One is that the survivorship function of one group—the "advantaged" group—never falls below that of the other (whether or not the two survivorship functions fully converge). The other is that both groups have an age by which all group members will become extinct, and the age for the advantaged group is not wildly older than that of the disadvantaged group (in most versions of the model, the two groups share an extinction age). These two assumptions jointly constrain the space of possible Gompertz parameters to ones in which the disadvantaged group has a larger intercept and the advantaged group has a larger slope of mortality over age. This combination of a larger intercept for one group and larger slope for the other produces a crossover in the hazards below the extinction point, as the advantaged group "catches up" in cumulative mortality. This model is precisely the one that we gave in Equation 3: a life course explanation of the crossover.

Liu and Witten's (1995) analysis is emphatically not a selection-based account, because the model contains no selection. That is, the crossover has nothing to do with compositional changes within each group: indeed, it on Liu and Witten's model, a crossover will occur even if the two subpopulations are each completely homogeneous, and mortality is stochastically determined solely by subpopulation membership and age. The crossover occurs not because the frailest members of the disadvantaged group (e.g., blacks) are selected out of the cohort, but rather

because an assumed pre-determined maximum lifespan for all individuals requires mortality to accelerate dramatically among all surviving individuals in the lower-mortality group.[8]

That Liu and Witten's model was an *alternative* to mortality selection, not an explication of it, was understood at the time (Olshansky 1995).[9] But in some recent work, it is prominently cited as an example of how mortality selection produces a crossover. For example, Dupre et al. (2006: 145) call it "a decidedly more rigid test of heterogeneity" than the mortality selection work that preceded it.[10] More startlingly, even the original proponents of the model sometimes misdescribe their own contribution. Hirsh et al. (2000:172) write that: "We have studied the relationships between populations which obey the programming scenario with respect to maximum lifespan and which, in consequence, display crossovers resulting from heterogeneity with respect to frailty (Liu & Witten, 1995)." This is simply not the case. The crossover between the advantaged and disadvantaged subpopulation in Liu and Witten's (1995) model does not derive from heterogeneity with respect to frailty because neither subpopulation is heterogeneous with respect to frailty. The model contains no heterogeneity, and therefore no selection.[11]

---

[8] Note that essentially all of the analytical work in producing the crossover is driven by the assumption that the groups have a such an extinction point, or maximum lifespan. (Liu and Witten suggested that the maximum lifespan was genetically determined.)

[9] In his introduction to the journal issue in which Liu and Witten's article appeared, Olshansky (1995: 583) correctly noted that the conclusions of the Liu and Witten model "would be expected to occur whether the populations being compared contain individuals who are homogeneous or heterogeneous in their endowment for longevity"—that is, their frailty.

[10] In some work, Liu and Witten (1995) has loomed even larger in discussions of the crossover. Johnson (2000) describes the Liu and Witten proposal as *the* crossover explanation, aside from the age misreporting hypothesis. It is unclear from her discussion, however, whether Liu and Witten's proposal is conceptualized as a mortality selection explanation or a life course explanation. In contrast to the emphasis placed on Liu and Witten's model by Dupre et al. (2006) and Johnson (2000), most crossover work has simply ignored Liu and Witten's proposal.

[11] In an about-face, Liu et al. (2008) develops a nearly identical model as the model in Liu and Witten (1995) and seems to present it as an alternative to a selection-based explanation of the crossover.

This problem of misdescribing life course explanations as mortality selection may have also afflicted one of the most prominent recent examinations of the crossover. In an article on religious participation and the crossover that acknowledges only mortality selection and age misreporting as possible explanations, Dupre et al. (2006) actually provide many of the elements of a life course explanation. They argue that religious participation lowers mortality, and that blacks are more often religious participants, which—since it lowers black mortality at the individual level—would seem to provide a straightforward life course hypothesis to account for the crossover. This account is only strengthened by Dupre et al.'s claim that religious participation furthermore lowers black mortality more than white mortality—a fact that would suggest that religious participants might reach a crossover at an earlier age than non-participants, which indeed is exactly what Dupre et al. find. Yet this hypothesis which arises so naturally from their premises and results is never discussed.

One aspect of Dupre et al.'s (2006) explication does seem to contradict a life course account: their hypothesis that the effects of religious participation on mortality will be greater earlier in the life course than at the old ages when the crossover is found. If this is true, it limits the extent to which a mortality-lowering effect of religious participation for blacks, more than for whites, could explain a crossover at old ages. However, the way this hypothesis is formulated is puzzling.[12] For one thing, if the crossover is explained by mortality selection, as Dupre et al.

[12] Here is the relevant passage in full:
"Second, to support the claim that religion contributes to why hazard rates invert, the effect of religion must vary with age. The reasons for this are twofold. First, heterogeneity in population frailty is generally greater at young ages. As populations age, mortality forces will disproportionately remove frail members from the population compared with those who are more robust. Therefore, the observed effect of attendance at religious services will be greatest earlier in the life course when the mortality-risk factors that attendance alters (e.g., unhealthy lifestyle) exert their strongest effects. Second, because the composition of the surviving population changes with age, favoring more-robust individuals, measures of heterogeneity will reflect this shift and exhibit lesser variation at older ages. Therefore, we hypothesize that the effect of religious involvement will be greater at younger ages (hypothesis H2b),

believe, then it is simply not true that, "to support the claim that religion contributes to why hazard rates invert, the effect of religion must vary with age" (Dupre et al. 2006: 146). An age-varying effect of religion on mortality is unnecessary whether the "effect" of religious participation is interpreted as an individual-level effect on mortality, or (what we believe the whole passage intends) the aggregate "effect" (e.g., the coefficient from a regression). For a mortality selection model to account for a mortality crossover, nothing at all needs to vary at the individual level, and the only thing that needs to vary at the aggregate level is how many members of each racial group are frail—no effect of anything on mortality needs to vary over age at all (other than the aggregate "effect" of being black, that is, the association between race and mortality).[13] Thus, the logical leaps in this passage may result in part from mixing together reasoning that applies to a mortality selection model (in which neither individual-level nor aggregate "effects" need vary with age) and reasoning that applies to a life course model (in which aggregate "effects" on mortality must vary with age, albeit in the opposite direction from the one hypothesized).

---

suggesting that reductions in mortality will favor those who experience higher mortality earlier in the life course. Together, the age- and race-dependent effects of attendance will produce an earlier crossover age as religious blacks gain a mortality advantage over religious whites before nonreligious blacks gain an advantage over nonreligious whites." (Dupre et al. 2006: 146)

It is hard to follow the reasoning in this passage because the points separated as "first" and "second" appear to say the same thing, and because neither sentence that begins "Therefore" seems to follow from the sentence preceding it without further unstated premises. However, the central claim made in the passage seems to shift between patterns of religious participation causing mortality selection ("to support the claim that religion contributes to why hazard rates invert"), and mortality selection causing patterns of religious participation (because the passage seems to argue that the mortality-protective benefits of religion accrue mostly to the frail, and hence are concentrated at earlier ages, when the frail are still alive).

[13] In the model given in Equations 1 and 2, no parameters vary with age, but mortality increases log-linearly with age for all individuals. In Vaupel and Yashin (1979), mortality selection (albeit in a single population, rather than a crossover) is illustrated in the case of mortality that is constant over age for the frail and robust (Vaupel and Yashin 1979: 176-177).

As a final note, we note that religious participation might produce a crossover even if the effect of religious participation on particular causes of death were proportionally constant—perhaps even if they declined over age—if the effect of religious participation varies across causes of death and the proportion of deaths due to each cause shifts over age. For example, imagine that religious participation has little effect on deaths to accidents, homicides, influenza, and colorectal cancers, but a large effect protecting against stroke, heart attack, and lung cancer (via encouraging healthy behaviors and reducing stress). Further imagine that the effect on this second set of causes of death is to reduce the risk of dying to these causes by a constant proportion over age. If the share of deaths shifts away from the first set of causes and toward the second set over age, then religious participation will be increasingly protective in absolute terms as a cohort ages even though its proportional effect on each cause of death has not changed.

Our point is not to conclude that religious participation instantiates a life course model of the crossover, but to elaborate some of the circumstances in which such a model might be plausible, and to illustrate the surprising absence of life course explanations of the crossover from any consideration even when they would seem to be clearly warranted. In sum, the failure to acknowledge life course-based explanations for the crossover as distinct from mortality selection has led analysts to incorrectly subsume them under the mortality selection's umbrella—to the detriment of our understanding of selection dynamics. If everything is described as selection, regardless of how it works, we will never be clear on how selection does work—or, more to the point, how it does not.

Conclusion

There are at least three ways that mortality crossovers can appear in data: data problems, selection at the aggregate level, and life course changes at the individual level. But when it comes to the black-white crossover, much recent demographic literature mentions only two explanations: racially differential age misreporting, and racially differential selective attrition by mortality. Drawing on a line of classic work within demography, we motivate renewed attention to the third possibility: that the crossover arises, in whole or in part, from a reduction of black disadvantage during the life course. We show that this life course theory lurks in the background in some important recent work on the crossover, even when it is not acknowledged explicitly.

We argue that ignoring the life course explanation, even in contexts that suggest its plausibility, has led some recent work to dramatically misdescribe the logic of mortality selection. Demography as a field is ill-served by verbally prioritizing mortality selection over competing accounts, if what the authors envision as the source of the crossover is not really selection after all.

The real stakes, in distinguishing mortality selection from life course changes experienced by individuals, are not merely theoretical, but substantive: these two explanations have quite different implications about the nature and extent of racial inequality at old ages, and accordingly, may bear on the etiology of racial inequality as well.

A pure selection account implies sustained black disadvantage that does not diminish over the life course. Applied to Fogel's (1986) data, this would be an understanding of slavery as a condition that disadvantages slaves compared to free citizens at every age. Applied to the modern United States, this would mean that neither the social conditions of old age (including

social programs targeted at the elderly, such as Medicare) nor the biological conditions of aging itself make a substantial dent in black disadvantage in survival at the individual level, whether because the sources of disadvantage (e.g., poverty, segregation, discrimination) are ongoing, or because disadvantages experienced earlier in life continue to reverberate at later ages.

By contrast, a pure life course story implies that black disadvantage reverses to white disadvantage at old ages. Applied to Fogel's data, slavery would be conceived (as it is in Fogel's [1986] work) as a condition that disadvantages the survival only of children, not adults. Applied to the modern United States, this explanation would imply that the life circumstances of elderly blacks are favorable to those of elderly whites, for example, because elderly whites are more likely to face the risks (including loneliness) of living alone.

Given the extent of black disadvantage across a diverse domains in the United States, few social scientists are likely to find such a pure life course story plausible. But many combinations rest in between these extremes: some factors might genuinely mitigate black disadvantage at old ages relative to younger ages (e.g., co-residence, Medicare, perhaps religiosity) and interact with heavier selection of blacks by mortality to produce the crossover. One important purpose in distinguishing such explanations is to be able to ask far more precise questions, such as: *how muc*h of the aggregate reduction in black disadvantage over the life course is "real" in the sense of reflecting declining disadvantage at the individual level, and how much of the aggregate reduction in disadvantage is due to selection? Given the potential interaction between the two explanations, what might mortality patterns look like if Medicare were scaled back, or alternatively, if anti-poverty programs targeting the elderly were expanded? Conversely, what might mortality patterns look like if child and adolescent mortality were made more equal

between blacks and whites in the United States, reducing the differential effects of selection? Such questions cannot be precisely posed and answered without engaging with how selection (in the aggregate) and life course changes (on the individual level) might each be implicated in the racial patterning of old-age mortality in the United States.

On the other hand, if we are unable to distinguish selection explanations from life course explanations, then mortality differentials at older ages can be used as evidence about inequality only given a prior assumption about how the evidence should be interpreted. Where Fogel saw only the life course, the recent literature on the crossover exclusively sees selection. With either assumption fixed firmly in place, mortality differentials then tell us primarily what we already assumed. This may not seem a hefty price to pay when one set of explanations conflicts with other social scientists' theoretical priors, as in Fogel's conclusion that slavery does not reduce adult survival. But in most contexts, this dependence of our interpretation of mortality disparities on mere assumptions about their most fundamental meaning—will a reversal of the hazards reflect a change in conditions at old age, or the lingering legacy of earlier inequality?—is indeed a large price. It is too large, we feel, for population science to bear. There are real sociological mechanisms that might account for black disadvantage in survival being smaller in the conditions of elderly Americans than in the conditions of younger Americans. There is a diverse range of crossovers, far beyond black and white, whose meaning ought to be determined by more than analytical preference. In the end, what is at stake in distinguishing—first theoretically, and ultimately empirically—the selection and life course explanations of the crossover is the need to rescue the status of mortality differentials as evidence, so that the conclusions we draw from it about the extent and nature of racial inequality are not driven solely by our prior beliefs.

REFERENCES

Berkman, L., Singer, B., & Manton, K. 1989. Black/White Differences in Health Status and Mortality Among the Elderly. *Demography*, 26(4): 661-678.

Coale, A.J., E.E. Kisker. 1986. Mortality Crossovers: Reality or Bad Data? *Population Studies* 40:389–401.

Corti, M.-C., Guralnik, J. M., Ferriucci, L., Izmirlian, G., Leveille, S., Paho, M., Cohen, H. J., Pieper, C., & Havlik, R. J., 1999. Evidence for a Black-White Crossover in All-Cause and Coronary Heart Disease Mortality in an Older Population: The North Carolina EPESE. *American Journal of Public Health*, 89(3): 308-314.

Jeptha P. Curtis, MD, Jared G. Selter, MD; Yongfei Wang, MS; Saif S. Rathore, MPH; Ion S. Jovin, MD; Farid Jadbabaie, MD; Mikhail Kosiborod, MD; Edward L. Portnay, MD; Seth I. Sokol, MD; Feras Bader, MD; Harlan M. Krumholz, MD, 2005. *Archives of Internal Medicine*, 165(1): 55-61.

Dupre, M. E., Franzese, A. T. & Parrado, E. A. 2006. Religious attendance and mortality: Implications for the black-white mortality crossover. *Demography,* 43**,** 141-164.

Dupre, M. E. 2007. Educational Differences in Age-Related Patterns of Disease: Reconsidering the Cumulative Disadvantage and Age-As-Leveler Hypotheses. *Journal of Health and Social Behavior*, 48(1): 1-15.

Elo, I.T. and S.H. Preston. 1994. Estimating African-American Mortality From Inaccurate Data. *Demography* 31:427–58.

Elwert, F. & Christakis, N. A. 2006. Widowhood and race. *American Sociological Review,* 71**,** 16-41.

Fogel, R. W. 1986. Nutrition and the Decline in Mortality Since 1700: Some Additional Preliminary Findings. *NBER Working Paper Series*. Working Paper No 1802.

Goldscheider, F. & Bures, R. 2003. The Racial Crossover in Family Complexity in the United States. *Demography*, 40(3): 569-587.

Guillot, M. 2007. "Mortality in Kyrgyzstan since 1958: Real Patterns and Data Artifacts." Espace Populations Societies 2007(1):113-26.

Hambright, T. Z. 1968. Comparability of Age on the Death Certificate and Matching Census Record, United States, May-August 1960. *Vital and Health Statistics Series*, Series 2, No. 9. National Center for Health Statistics.

Hill, M. E, Preston, S. H., and Rosenwaik, I. 2000. "Age Reporting Among White Americans Aged 85+: Results of A Record Linkage Study." Demography 37(2): 175-86.

Hirsch, H. R., Liu, X., and Witten, T. M. 2000. Mortality-rate crossovers and maximum lifespan in advantaged and disadvantaged populations: Accelerated-mortality and sudden-death models. *Journal of Theoretical Biology*, 205(2): 171-180.

Hoffman, F.L. 1896. The race traits and tendencies of the American Negro, *Publications of the American Economic Association* 11: 1–329.

Hoffmann, R. 2008. *Socioeconomic Differences in Old Age Mortality*. Springer: New York, NY.

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.

Huang, F., and Wu, C. 2010. "Old Age Mortality Crossover in China: Does the Public Health Insurance Matter?" *Population Research* 34(1):95-105 [in Chinese; abstract in English].

Johnson, N. E. 2000. The racial crossover in comorbidity, disability, and mortality. *Demography*, 37(3): 267-283.

Kestenbaum, Bert. 1992. "A Description of the Extreme Aged Population Based on Improved Medicare Enrollment Data." Demography 29(4): 565-80.

Li, T., Yang, Y., & Anderson, J. 2013. Mortality Increase in Late-Middle and Early-Old Age: Heterogeneity in Death Processes as a New Explanation. *Demography*, 50(5): 1563-1591.

Liu, X., Engel, C. C., Armstrong, D. W., and Kang, H. 2008. Survival convergence and the preceding mortality crossover for two population subgroups. *Population Research and Policy Review* 27(3): 293-306.

Liu, X. and Witten, M. 1995. A biologically-based explanation for mortality crossovers in human populations. *Gerontologist*, 35(5): 609-615.

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.

Manton, K. G., Poss, S. S., & Wing. S. 1979. The Black/White Mortality Crossover: Investigation from the Components of the Perspective of Aging. *The Gerontologist*, 19(3): 291-300.

Manton, K. G. & Stallard, E. 1981. Methods for evaluating the heterogeneity of aging processes in human populations using vital statistics data: Explaining the black-white mortality crossover by a model of mortality selection. *Human Biology,* 53**,** 47-67.

Masters, R. K. 2013. Uncrossing the U.S. Black-White Mortality Crossover: The Role of Cohort Forces in Life Course Mortality Risk. *Demography*, 49(3): 773-796.

Nam, C. 1995. Another look at mortality crossovers. *Biodemography and Social Biology*, 42:1-2, 133-142.

Olshansky, S.J. 1995. New Developments in Mortality. *The Gerontologist* 35(5):583-587.

Parnell, A. M. & Owens, C. R. Evaluation of U.S. mortality patterns at old ages using the Medicare Enrollment Data Base. *Demographic Research*, 1, 2.

Pearl, Raymond. 1922. The Biology of Death. J.B. Lippincott: Philadelphia, PA.

Preston, S.H., Elo, I. T., Rosenwaike, I. and Hill., M. E. 1996. African-American Mortality at Older Ages: Results of a Matching Study. *Demography* 33:193–209.

Rosenwaike, Ira. 1981. A Note on New Estimates of the Mortality of the Extreme Aged. *Demography* 18(2): 257-66.

Ruggles S. R. 1994. The Origins of African-American Family Structure. *American Sociological Review,* 59: 136-151.

Sagi-Schwartz, A., Bakermans-Kranenburg, M. J., Linn, S., & van IJzendoorn, M. H. 2013. Against All Odds: Genocidal Trauma Is Associated with Longer Life-Expectancy of the Survivors. *Plos One*, DOI: 10.1371/journal.pone.0069179

Sautter, J. M., Thomas, P. A., Dupre, M. E., & George, L. K. 2012. Socioeconomic status and the Black-White mortality crossover. *American Journal of Public Health*, 102(8): 1566-71.

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.

Wilmoth, J. R. & Dennis, M. 2007. Social Differences in Older Adult Mortality in the United States: Questions, Data, Methods, and Results. In Robine, J. M., Crimmins, E. I., Horiuchi, S., & Zeng, Y. *Human Longevity, Individual Life Duration, and the Growth of the Oldest-Old Population*. International Studies in Population.

Yi, Z., & Vaupel, J. W. Oldest-Old Mortality in China. *Demographic Research*, 8(7): 215-244.

CHAPTER 4:

The Selection Problem for Testing the Selection Explanation of the Black-White Mortality Crossover

Introduction

Mortality crossovers abound: it frequently occurs that one population has higher mortality than another at young ages, but lower mortality at older ages. Crossovers have been found between Navajo Native Americans and the total U.S. population (Thornton 2004), between smokers and non-smokers in the United States (Hoffman 2008), between Russian and Kyrgyz infants in Kyrgyzstan (Guillot 2007), between individuals with public health insurance and those without in China (Huang and Wu 2010), and in national mortality between Sweden, Japan, and the Han Chinese of China (Yi and Vaupel 2003).

By far the best-known, and best-documented, mortality crossover is that between African Americans and whites in the United States. Hoffman (1896) and Sibley (1930) provide the earliest reports of crossovers in black-white age-specific death rates at old ages. Despite concerns of measurement problems in some datasets, recent corrections for measurement error do not fully eliminate the black-white mortality crossover (Hill et al. 2000), and other recent work documents black-white mortality crossovers in various high quality datasets with rigorous age reporting, such as Medicare and EPESE (e.g. Kestenbaum 1992, Dupre et al 2006, Sauter et al. 2012). The black-white mortality crossover is "real." The question is what it tells us about heterogeneity within black and white cohorts and inequality between them.

The black-white mortality crossover is usually assumed to reflect mortality selection: black and white cohorts are each heterogeneous in the risk of dying, which produces systematic

selection for the most mortality-resistant individuals as they age. Because black mortality is higher, black cohorts are selected more sharply: if black and white cohorts start out with the same proportion of 'frail' and 'robust' members at birth, then at every subsequent age the proportion frail among the survivors will be smaller in the black cohort than the white one.[1]

That mortality selection might produce a crossover has been understood for at least 35 years (e.g., Vaupel et al. 1979, Vaupel and Yashin 1985). What has lagged behind has been a more detailed understanding of the substance of selection: what is the *frailty* that is selected out of cohorts, faster for blacks than for whites, and how, precisely, does this selection proceed as the cohorts age? Over the same period of time that the crossover has become a widely understood phenomenon in demography, the demography of health more generally has exploited (and collected) more detailed and diverse data sources to delve into the specific biological and social mechanisms that produce health inequalities. In recent years, this general turn has affected research on the crossover as well. Demographers have begun to put socioeconomic substance onto the structure of mortality selection models by conditioning mortality on particular dimensions of heterogeneity—such as income and education (Sautter et al. 2012), neighborhood (Yao and Robert 2011), and even religious participation (Dupre et al. 2006)—and analyzing the age at crossover in the more specific subpopulations that result.

This paper evaluates when the age at crossover, conditional on some form of heterogeneity, is informative about an underlying mortality selection model. In principle, the age at crossover might be informative in one of two ways. We might wonder whether selection accounts for the crossover at all: if we believed that a particular dimension of heterogeneity

[1] This claim is true as long as the frail and robust have the same mortality ratio in the black and white cohorts, but with a higher level of each in the black cohort.

being selected out of one cohort faster than another produces the crossover, but find that mortality analyses refute this hypothesis, then perhaps we will conclude that the crossover reflects, instead of selection, a genuine reversal of disadvantage over the life course. Or we might feel confident in assuming that mortality selection is the ultimate cause of a particular crossover, but wonder what role in that process is played by a particular dimension of heterogeneity—that is, inequality—such as income disparities. In either case, what we need is a selection model that makes testable predictions about mortality patterns. If such predictions are violated, then the analyst must decide whether the conclusions should be local in scope (the particular heterogeneity investigated does not produce the crossover) or global (mortality selection in general is not the cause of the crossover). But a precondition for any conclusion is a real prediction from a selection model. Yet although the strategy of exploring the age at crossover, conditional on social heterogeneity, is increasingly prominent in the crossover literature, what conclusions are licensed by this strategy remains unclear. We investigate what conditioning on isolated dimensions of heterogeneity, and comparing the resulting ages at crossover, may or may not reveal about inequality between blacks and whites.

Our approach seeks to build on the strengths of what we characterize as two distinct traditions in the wide-ranging crossover literature. The ‘classical period’ of work on mortality selection (roughly, the 1980s and 1990s) produced work that was in many ways theoretically richer than it is remembered as being; in particular, in this period, work on mortality selection often sought to build precise models of the sometimes surprising ways that individual-level patterns of inequality and status transitions can be reflected in population aggregates.

More recently, from roughly 2000 onward, this approach seems to have been largely abandoned. Little of the recent work on the crossover commits to any specific selection model at all. At the same time, however, work on the crossover has usefully broadened its substantive understanding of heterogeneity, notably beginning to incorporate dimensions of social inequality into work on mortality selection. Unfortunately, because this broadening has happened without the benefits of the earlier attention to precisely relate individual trajectories to changing population compositions, the theoretical significance of this work has sometimes been obscured and inaccurately described.

We argue that these two strands of work on the crossover must be melded. By broadening the substantive content of heterogeneity to focus centrally on socioeconomic inequalities, the crossover literature has begun to *implicitly* embrace models of multidimensional, dynamic heterogeneity that are far more substantively realistic than the models that first fixed demographic attention on mortality selection. Yet by abandoning the demographic tradition of building explicit population models, this recent crossover literature has sometimes drawn conclusions that are not well-supported. Unfortunately, we show that when the selection dynamics of multidimensional, dynamic heterogeneity are accurately described, they become essentially untestable using the kinds of statistical adjustments pursued in recent work on the crossover.

Mortality Selection and Heterogeneity: The Classic (Unidimensional) Model

This section presents an explicit mortality selection model based on fixed, unidimensional heterogeneity, which serves as a baseline for the multidimensional heterogeneity model

developed below. Imagine a cohort consisting of four groups with fixed membership until they die: frail and robust blacks and whites. Let each group have Gompertz mortality and let the frail and blacks each raise mortality proportionally over age. The resulting age-specific mortality of each subgroup is given in Equations 1(a-d), with subscripts identifying the groups as white or black and robust or frail. In these equations, mortality at a given age *a*, μ(a), depends on the rate of exponential mortality increase over age, β, which is assumed to be equal for all groups; on an intercept value (mortality at baseline age) for robust whites, α; a mortality multiplier for the frail, *f*, which is assumed to be equal for blacks and whites and assumed to exceed 1; and a mortality multiplier for blacks, *b*, which is assumed to be equal for the frail and the robust and assumed to exceed 1.[2]

$$\begin{aligned}
\mu_{w,r}(a) &= \alpha e^{\beta a} \\
\mu_{w,f}(a) &= f\alpha e^{\beta a} \\
\mu_{b,r}(a) &= b\alpha e^{\beta a} \\
\mu_{b,f}(a) &= bf\alpha e^{\beta a}
\end{aligned} \tag{1}$$

Within each group, black mortality exceeds white mortality at each age: there is no crossover at the level of individuals. Any crossover in the aggregate will come from changes in population

[2] These are standard assumptions in the mortality selection literature (although occasionally relaxed; see Horiuchi and Wilmoth (1997) and Lynch et al. (2003) for discussions of heterogeneous slopes over aging, vs. the proportional mortality differences considered here). Our discussion will retain most of these premises in order to focus on the consequences, for selection dynamics, of some particular aspects of the *frailty* construct, as explained below.

composition. The age-specific mortalities of blacks and whites, aggregated over frailty, are given in Equation 2. This race-specific aggregate mortality is an average of the frail and robust groups, weighted by their proportion in the cohort at the given age. The proportion frail is represented by π(a) for each race.

$$\bar{\mu}_k(a) = \pi_k(a)\mu_{k,f}(a) + \big(1 - \pi(a)\big)\mu_{k,r}(a) \tag{2}$$

The key dynamic of mortality selection is that the proportion frail declines over age, shifting the weight of aggregate mortality toward the robust, and that the proportion frail declines faster among blacks. The age-specific proportion frail for each racial group is given in Equation 3a. In Equation 3a, α* is used as a shorthand representing the intercept (mortality at baseline age) for each group. Thus, α* is equal to α for whites and α*b* for blacks, the white intercept times the black mortality multiplier; the equation is otherwise identical for the black and white groups. The other new term in Equation 3a is *p*, which gives the ratio of frail to robust members of the racial group at baseline; hence, we assume that blacks and whites had the same proportion frail from baseline (e.g., from birth), in order to focus on the subsequent selection dynamics.

$$\pi_k(a) = \frac{p}{p + \big(\exp(\exp(\beta a) - 1)\big)^{(f-1)\frac{\alpha^*}{\beta}}} \tag{3a}$$

The expression $\big(\exp(\exp(\beta a) - 1)\big)^{(f-1)\frac{\alpha^*}{\beta}} = \frac{S_{kr}}{S_{kf}}$ equals the ratio of the survivorship function (i.e., the proportion surviving) of the robust to the survivorship function of the frail. This lets us rewrite the proportion frail in more intuitive terms, in Equation 3b:

$$\pi_k(a) = \frac{p}{p + \frac{S_{kr}}{S_{kf}}} = \frac{p}{p + \omega_k} \tag{3b}$$

Since the ratio of the survivorship of the robust to the survivorship of the frail will recur in later equations, we christen it ω_k. The properties of the formula in Equation 3b for the proportion frail among the survivors to each age are sensible and intuitive: since all variables are positive, the proportion frail is always non-negative and always less than 1; a larger ratio of frail to robust at baseline produces a larger proportion frail at subsequent ages, all else equal (since *f* is inversely related to ω_k); and the larger the frailty multiplier, *f*, the smaller the surviving proportion frail at a given age.

One property plays a special role in the understanding the crossover: the larger the mortality intercept, α*, the fewer frail survive to a given age. This is the key fact that underpins how mortality selection might make black mortality drop below white: since black mortality exceeds white mortality in individuals, represented here with the black multiplier in the α* term, the black group comes to have a lower proportion of surviving frail members.[3]

Given these equations, we can derive the conditions for a crossover to occur in the unidimensional heterogeneity model. Equation 4 gives the black-white aggregate mortality difference at each age.

[3] This conclusion does depend on the model form. Specifically, the conclusion follows from the assumption (widespread in mortality studies, not only in studies of the crossover) that covariates have a proportional effect on mortality. This allows the difference between frail and robust mortality, also assumed to be proportional to one another, to be larger in absolute terms for groups with generally higher mortality. That absolute difference between frail and robust mortality determines, in part, the rate of change in the frailty of the survivors. (For more discussion of the rate of selection, see Vaupel and Zhang (2010) for the general case and Wrigley-Field (2014) for the case of binary frailty.)

$$\bar{\mu}_b(a) - \bar{\mu}_w(a) = \alpha e^{\beta a}\left[\left(f-1\right)\left(b\pi_b(a) - \pi_w(a)\right) + b - 1\right] \tag{4}$$

In Eq. 4, the terms $\alpha e^{\beta a}$ (i.e., the mortality of robust whites at age *a*), $b-1$ (i.e., the extent to which black mortality is elevated above white mortality at the individual level), and $f-1$ (i.e., the extent to which frail mortality is elevated above robust mortality) are always positive. The only term that can be negative, and therefore can lead to the black-white mortality crossover, is the term $b\pi_b(a) - \pi_w(a)$. Thus, Equation 5 derives from Equation 4 the conditions in which black mortality will fall below white mortality:

$$\bar{\mu}_b(a) < \bar{\mu}_w(a) \leftrightarrow \left(\pi_w(a) - b\pi_b(a) > \frac{b-1}{f-1}\right) \tag{5}$$

In other words, the black and white hazards cross if and only if the difference between the proportion frail among white survivors and the proportion frail among black survivors times the black mortality multiplier is larger than the ratio between the extent to which black mortality exceeds the mortality of whites with the same frailty and the extent to which frail mortality exceeds the mortality of the robust of the same race. One intuitive consequence: there is no crossover if being black raises mortality more than being frail at the individual level (since the left hand side can never equal or exceed 1). A crossover, then, is more likely to happen when the proportion frail among white survivors is much larger than the proportion frail among black survivors, when the individual-level mortality difference between blacks and whites is relatively

low, and when that difference is greatly exceeded by the difference between frail and robust mortality. In what follows, Equation 5 will provide a key link between the conditions for a crossover in the unidimensional heterogeneity model and in the multidimensional heterogeneity model.

This model represents the classic mortality selection explanation of the crossover.[4] Recent work on mortality selection relies less directly on theoretical constructs like frailty, and instead uses particular traits—particular components of the heterogeneity within cohorts, such as socioeconomic status—in an attempt to elucidate more specific selection dynamics.[5] Yet the relationship between measurable traits that may account for the crossover, and the *frailty* construct driving mortality selection, is not sufficiently considered in this work. As a first pass, we can say that a trait is part of frailty—that it is the kind of heterogeneity that, by being selected out of a cohort, can effect a mortality crossover—if it is fixed in individuals *and* has constant sign in its relationship to mortality.[6] Genes are largely fixed in individuals, though their

---

[4] Often, however, frailty is assumed to follow a Gamma distribution (e.g., Gampe 2010; Horiuchi and Wilmoth 1998; Missov and Finkelstein 2011; Vaupel et al. 1979), rather than the binary frailty used in this model. For a discussion of Gamma-distributed vs. binary frailty, see Wrigley-Field (2014:66-67). We use binary frailty in this paper because it simplifies the discussion of the complex dynamics of multidimensional heterogeneity, allowing for greater intuition into how the multiple dimensions are made, by mortality selection, to interact.

[5] Theoretical discussions of mortality selection usually speak of *frailty* (vs. *robustness*), whereas recent empirical work often relies more heavily on the term *heterogeneity*. While the latter term has broader connotations in everyday usage, the two terms play the same theoretical role in mortality selection theory, and so we use them interchangeably. In the bulk of the discussion below, we refer to *heterogeneity* simply because *frailty* is more likely to be misread as implying a more narrowly biological construct than is intended. (For example, it may feel wrong to some readers to say that people who never complete high school are *frail*, but hopefully it does not strike any readers as wrong to say that they are less likely than others in the United States to survive to old age. These formulations denote, but to many people do not connote, the same claim.)

An alternative, and quite defensible, terminological choice would be to use *heterogeneity* for the broad set of differences among individuals within a racial group, while reserving *frailty* for those differences that do behave like the frailty of mortality selection theory, e.g., one could reserve *frailty* for factors that together raise individual-level mortality by a constant proportion across age. This usage has the advantage of allowing finer conceptual distinction, but we fear it has the disadvantage of lending itself too easily to preconceptions about the substantive nature of frailty.

expression is not; but some social advantages and disadvantages are also fixed at birth (such as one's grandparents' education), and others are quite stable over an individual lifespan. Childhood exposure to lead, a criminal record or dishonorable discharge earned in one's youth, a lack of documentation of one's immigration—all might, in principle, raise the risk of dying for the same individuals at age after age, and so be the kind of thing culled from cohorts as they grow older. Such *durable inequalities* are what mortality selection theory is about.

But in fact, as we will show, moving from abstract frailty to particular traits is not so straightforward. Between the theoretical elegance of *frailty* and the empirical specificity of particular, measurable traits lies a cavern whose contours remain unexplored. The aim of this article is to shine a light on some of those contours, to place the first foundations of a bridge between the theory (mortality selection) and the practice (adjusting for "heterogeneity") of research into how mortality selection might make a crossover.

Testing Selection Models with Age at Crossover

An increasingly common and prominent approach to the crossover is to examine the age at crossover conditional on a single dimension of heterogeneity. For example, Dupre et al. (2006) find that, among women, the crossover occurs at age 80 among women who regularly participate in religious services, age 90 among women who do not, and age 85 among women aggregated over religious participation. Sautter et al. (2012) find a crossover occurring among low-income men before age 65 (the youngest age in their data); among high-income men, at an age older than 65 (the specific age is not reported); and at age 79 among men in the aggregate. Yao and Robert

[6] See Wrigley-Field and Elwert (n.d.) for elaboration of the conditions in which mortality selection will occur.

(2011) find a crossover at age 74 when controlling for individual and neighborhood socioeconomic status, compared to age 80 in the population as a whole.

By putting social heterogeneity at the center of the crossover literature, works employing this strategy have led a shift away from idealized frailty models and toward substantive engagement with stratification processes that may underlie selection models. However, this new crossover literature has moved away from the classical crossover tradition in more than its expanded interpretation of heterogeneity. While implicitly drawing on more complex selection models than have traditionally been invoked, this literature has also ceased to formulate explicit models of the selection processes that are being evaluated or assumed. The result is inferences about selection that are at times ambiguous, and at other times unambiguously mistaken.

In general, the purpose of estimating the age at crossover has been unclear because it is not clear what question is being addressed: is the purpose of the analysis to test mortality selection theory (against what alternative?), or to assume that mortality selection explains the crossover but test whether the particular trait explored is (part of) the "frailty" that selection acts on, or something different?

Where the analysis is explicitly linked to conclusions about inequality, the conclusions do not always follow. For example, the conclusions that Sautter et al. (2012) draw from their analysis misconstrues the basic fact about mortality selection, namely that, all else equal, mortality selection causes a population with higher early-life mortality to have lower later-life mortality. Sautter et al. (2012: 1566) summarize the implication of their investigation of mortality selection as: "A primary route to reducing mortality differentials in later life is to prevent the disproportionate selective mortality of Blacks and the poor earlier in the life

course."[7] But this is the opposite of what mortality selection theory predicts. Reducing the disproportionate selective mortality of blacks and the poor would increase, not reduce, black and poor mortality relative to white and non-poor mortality at older ages. This group-level relationship between higher early-life mortality and lower later-life mortality is the fundamental point about mortality selection, and the theoretical underpinning of the mortality selection explanation of the crossover.

In other ways, the basic logic of the prominent recent analyses is not clearly explicated. Sometimes the main comparison made is the age at crossover between those who have the explored trait and those who do not (e.g., Dupre et al. 2006); sometimes it is the age at crossover between the aggregate population and the population conditioned on the trait (e.g., Sautter et al. 2010, Yao and Robert 2011). In neither case is it clear what crossover patterns would license what conclusions. In particular, two questions loom over the literature that conditions the age at crossover on measured traits: it is unclear what role in a mortality selection model the measured trait is meant to be playing, and it is unclear whether particular subgroups should be expected to cross at *older* or *younger* ages than whatever group they are compared to.

Regarding the first question, for example, in Dupre et al. (2006), it sometimes seems as though the point is that religiosity *is* the heterogeneity that selection acts on,[8] and other times as

[7] Similarly: "These findings suggest that, although public health efforts are needed during later life to provide income adequacy to all older adults, the primary route to reducing mortality differentials in later life is to prevent the disproportionate selective mortality experienced by Blacks and the poor earlier in the life course." [Sautter et al. 2012: 1570] This may be true, but surely the crossover is not evidence of it: were we to prevent disproportionate selective mortality in blacks and the poor at young ages, mortality selection theory would tell us to expect *greater* mortality differentials, in the direction of those groups' apparent disadvantage, at older ages.

[8] For example: "Yet, a noticeable limitation of much of the current research on heterogeneity is the reliance on post facto explanations used to account for the existence of a crossover. Only a few studies have directly incorporated any biological, social, or environmental covariates to help isolate the sources of heterogeneity that may produce a crossover. […] With the recent attention given to the relationship between religion and mortality, we extend this line

though the point is that religiosity intensifies or lessens the selection on some *other* heterogeneity that drives the crossover.[9] But these are conflated, rather than clearly distinguished as two separate roles that religiosity might play. Likewise, Sautter et al.'s (2012) various formulations potentially offer *four* distinct relationships that income might have to the crossover. Analogously to Dupre et al., it is generally unclear whether income is itself the "heterogeneity" whose racially differential selection by mortality produces a crossover, or whether low income contributes to a crossover by intensifying black selection for robustness earlier in life.[10] Inversely, other times it sounds as though the selective mortality of frail blacks is supposed to explain the association between income and mortality, instead of the other way around.[11] And still other times, it sounds

---

of work and ask whether religious attendance plays a role in a black-white mortality crossover." (Dupre et al. 2006: 145)

[9] For example: "The following research questions focus on whether the black-white mortality crossover is also associated with religious attendance. The assumption of this research is that a crossover occurs largely because of heterogeneity: that is, a crossover occurs because some groups in a population have a greater susceptibility of dying and systematically experience higher mortality at younger ages, creating compositional changes among survivors. As a result, the population is increasingly composed of robust black survivors and a mixture of frail and robust whites, who eventually exhibit mortality rates that are higher than blacks. *We hypothesize that religious involvement affects the mortality crossover by differentially influencing the mortality rates of blacks, particularly at younger ages. In other words, decreases in black mortality at younger ages, vis-a-vis religion, increases the extent of heterogeneity in population frailty among blacks and whites at older ages*." (Dupre et al. 2006: 146; emphasis added)

[10] The formulations in Sautter et al. (2012) article repeatedly refer to income as a "source of heterogeneity in frailty," and it is not clear to us from the analysis which of these two interpretations (i.e., income is itself the heterogeneity, or low income increases the extent of heterogeneity in some other frailty) is intended.

[For example: "Only a handful of studies have incorporated substantive covariates to help identify sources of heterogeneity in frailty. We hypothesized that the unequal distribution of socioeconomic status (SES) contributes to heterogeneity in frailty and the mortality crossover paradox." (Sautter et al. 2012: 1566); also: "Our study was the first to our knowledge to examine whether low education and income capture important sources of heterogeneity in individual frailty that contribute to the Black-White mortality crossover." (Sautter et al. 2012: 1566) Similarly: "Third, low income is associated with increased mortality for both men and women and is an important source of heterogeneity in frailty that has an impact on the Black-White crossover for men, but not women." (Sautter et al. 2012: 1568)]

No more specific formulation is offered to lend weight to either of these readings over the other, except for the quotes provided in the next three footnotes, which in our reading support still other models than either income as frailty or as a mitigator of selection against frailty.

[11] For example: "The finding that low income becomes less important for Black men than for White men over time is interesting and may reflect 2 processes. First, because selective mortality has already taken a greater toll on Black

like the mortality-raising effect of having a low income changes over age at the individual level, contributing to a crossover[12] (which, although it is discussed as part of a mortality selection theory, is actually an alternative explanation to mortality selection—an example of what we have elsewhere [in Chapter 3] called the life course alternative). None of these four possibilities is attached to any clear model or prediction about mortality selection theory, and in general it is not made clear that they are distinct possibilities or how they might interact. We agree that all four may be possible, but to subsume them all under labels like "explain[ing] heterogeneity in frailty" (Sautter et al. 2012: 1569) or "captur[ing] important sources of heterogeneity in frailty that contribute to the black-white mortality crossover" (Sautter et al. 2012: 1566) provides little basis for either understanding or testing mortality selection.

Predictions about the order in which groups should reach a crossover remain equally unclear for analysts hoping to build on the existing literature. For example, Dupre et al. (2006) shows that the religious, the *lower*-mortality group, cross at a younger age than the non-religious, while Sautter et al. (2012) show that the poor, the *higher*-mortality group, cross as a younger age than the non-poor; it is not clear what predictions would be licensed in future research. Equally, it is not always clear how the observed changes in crossover timing relate to mortality selection: Sautter et al. (2012) conclude that socioeconomic status "explains" the crossover because they find that adjusting for income moves the crossover to *younger* ages, which on its face appears to

---

men, the more robust survivors may be less sensitive to economic disadvantage than their White counterparts may be." (Sautter et al. 2012: 1569) (The second hypothesized process referred to is the individual-level argument about Medicare quoted in the next footnote.)

[12] For example: "Income explained the Black-White mortality crossover presumably because older Black men faced reduced risk associated with low income and this risk further attenuated with increasing age." (Sautter et al. 2012: 1569); also: "Medicare may have a larger impact for Black men than for White men because it ameliorates some problems associated with economic deprivation and Black men were disproportionately more likely than White men to lack health insurance before the age of 65 years." (Sautter et al. 2012: 1569)

be not explaining the crossover, but exacerbating it. Since they provide no explicit criteria for what would constitute evidence for or against a trait explaining the crossover,[13] it is unclear whether, if adjusting for income had moved the crossover to older ages rather than younger ones, any of the substantive conclusions would change or not.

Thus, in general, the recent empirical literature as a whole provides little guidance to those hoping to use the age at crossover, conditional on some social heterogeneity, to draw substantive conclusions about processes of inequality and changing cohort composition over the life course. We argue that these weaknesses have a common source: *The overarching problem with the way this literature has proceeded is that there has been no attempt to grapple with the gap between the theoretical construct of frailty and the practice of adjusting for particular socioeconomic inequalities.*

It is clear that what analysts have in mind is a more complex, more richly realistic heterogeneity model than the stylized frailty models evoked in the traditional literature (e.g., Manton and Stallard 1981, Vaupel et al. 1979, Vaupel and Yashin 1985), such as the model given in Equations 1-5. After all, if income, neighborhood, or religious participation are dimensions of the kind of heterogeneity that can propel mortality selection, then the kind of

---

[13] Sautter et al. (2012) conclude that a trait is *not* part of the heterogeneity that explains the crossover when regression of mortality has no statistically significant interaction term between the trait and race.

(They say: "The main effect of low income was associated with increased mortality risk for both men and women, but did not vary by race for women. Thus, although inclusion of income and its interactions rendered the parameters of the Black-White mortality crossover (i.e., race and race a age) nonsignificant, income was not a source of heterogeneity in frailty that had an impact on the Black-White mortality crossover for older women. Among women, it is possible that social resources are more strongly associated with heterogeneity in frailty than are socioeconomic resources." (Sautter et al. 2012: 1569-1570))

However, it is not clear what their positive criterion is for a trait to be a "source of heterogeneity in frailty" affecting the crossover.

(In fact, such an interaction effect is neither necessary nor sufficient for a trait to act like 'frailty' in the standard mortality selection model considered here, since the action in a mortality selection model is not in the effect sizes, but in the group composition. Indeed, that no race-frailty interaction is needed to produce differential selection by race is in some sense the main point of the classic crossover model [e.g., Vaupel et al. 1979].)

heterogeneity that propels mortality selection simply does not work like the unidimensional fixed heterogeneity of mortality selection theory. Since presumably no one thinks that neighborhood, for example, captures the *entirety* of this heterogeneity, the heterogeneity must be multidimensional; and traits like one's neighborhood are gained and lost over the life course, so the heterogeneity must be dynamic.

Yet none of this work engages with the difference between these implicit conceptions of heterogeneity and how that difference changes the predictions one can draw based on mortality selection theory—and hence, ultimately, the inferences we can draw about mortality selection based on conditioning on these traits. The impulse seems to be something like: we understand from a well-developed theoretical literature how black-box *frailty* works; if we pull out particular traits that help to comprise that frailty, they will work in a similarly understandable way. Unfortunately, this assumption turns out to be wrong.

In the remainder of this article, we engage with this gap between the explicit conception of heterogeneity in much of the theoretical literature on mortality selection and the implicit conception in the recent empirical literature, and we show that the implications of the change are profound. In light of the ambiguities in the literature, we will develop a framework that: clarifies each of the roles in a selection model that a single dimension of heterogeneity might play, and how they interact; uses the resulting selection model to derive predictions about when subgroups, conditional on some trait, might cross at older or younger ages than each other and the aggregate, as well as clarifying when such predictions are not well supported; and uses these predictions to discuss the circumstances in which conditional on a particular dimension of heterogeneity can serve as a test of a mortality selection model. Such a formal analysis is needed to rejoin the

proximate project of conditioning on heterogeneity and assessing the age at crossover to the ultimate project of using mortality to understand inequality.

Unidimensional Fixed Heterogeneity

Unidimensional fixed heterogeneity is the canonical case. It is the model we gave in Eqs. 1-5, reflecting 35 years of research into mortality selection. For concreteness, assign to the frailty term in that model, *f*, a particular trait: imagine it is the trait of having been exposed to tobacco as a fetus.[14] This trait is clearly fixed by the time of birth. We can further imagine that it raises mortality at all ages, and that it is equally distributed at birth between blacks and whites.[15] According to mortality selection theory, the tobacco-exposed tend to die young in both black and white cohorts, but the combined effects of fetal tobacco exposure and the disadvantages of being African-American combine to make the black cohort become disproportionately composed of non-exposed more quickly than the white cohort.

What happens when mortality is stratified on fetal tobacco exposure (i.e., frailty)? If selection by mortality against fetal tobacco exposure accounts for the crossover, then the black and white hazards among those who were exposed, and those who were not exposed, should have no crossovers. Since within each group, blacks and whites have identical frailty, black

[14] Throughout the discussion in this article, we use binary examples of heterogeneity in order to focus the attention on the complications that are central to our argument—multidimensionality and dynamism—by simplifying others. But as the example highlights, such binary traits are often simplifying a continuous reality: a fetus is not just exposed or not exposed to tobacco, but can be exposed in varying amounts, on various schedules, and at various developmental moments. (For a general discussion of the substantive meaning of binary vs. continuous frailty in mortality selection models, see Wrigley-Field 2014: 66-67.)

[15] This assumption is not meant to be realistic; this is a stylized example to illustrate the logic of the selection model. The rates of smoking among black and white women are similar (Ho and Elo 2013), but the rates of prenatal exposure depend on the racial patterning of the probability of smoking conditional on pregnancy; the correlation of this probability with family size (to translate women's probabilities into children's (Preston 1976)); and fetal survival to birth, conditional on smoking.

mortality should dominate white mortality at all ages within the exposed and the non-exposed, even though in the aggregate, black mortality falls below white mortality.

Of course, it is absurd to imagine that fetal tobacco exposure is the only kind of heterogeneity that contributes to the stable ordering of individual mortality over age. By the same token, nobody who empirically analyzes income, neighborhood, or religiosity as a particular dimension of heterogeneity believes that it fully accounts for the crossover, the way it does in idealized frailty models. But a naive extension of unidimensional heterogeneity would expect statistical controls for fetal tobacco exposure to still behave much as statistical controls for *f* do, only less completely. Thus, if the single dimension of heterogeneity does not fully account for the crossover—if conditioning on it does not actually remove the crossover from each group—then perhaps it partially accounts for it, by moving the crossover to an older age than it occurs at in the aggregated cohort, so that controlling for more and more heterogeneity would move the crossover to older and older years, until it disappeared. If fetal tobacco exposure is a *part* of the heterogeneity that gets selected out of the black cohort, the reasoning would go, then adjusting for that single component should make the composition of black and white survivors closer to equal, if not fully equal. And if those compositional differences, of which fetal exposure is a component, are what makes black mortality fall below white mortality, then—the reasoning goes—adjusting for the single component, fetal exposure, should raise black mortality relative to white. It should push the crossover to an older age.

It would naively appear, then, that Sautter et al. (2012) have disconfirmed the selection story that uses poverty as a type of heterogeneity producing the crossover. Not so fast: this appearance is misleading. The error is in presuming that the rich and complex heterogeneity

implicitly invoked in empirical work (work that, for example, conditions on poverty) behaves like the unidimensional heterogeneity of theoretical work. Conditioning on unidimensional heterogeneity removes the crossover. But in multidimensional heterogeneity, the dimensions interact so that each dimension is not simply a lesser version of this unidimensional model—the crossover may not disappear, but it surely will happen later—but instead behaves in ways that are qualitatively more complex. We now turn to the dynamics of multidimensional heterogeneity.

Multidimensional Fixed Heterogeneity

To see why the naive extension of the unidimensional model does not work, consider two traits. As a similarly stylized example, imagine that, in addition to fetal tobacco exposure, there is a genetic predisposition to cancer, equally distributed between blacks and whites at birth, that behaves just like fetal tobacco exposure did in the previous example, is distributed independently of it, and raises mortality independent of it. Imagine that fetal tobacco exposure and this cancer gene together constitute the entirety of the heterogeneity that produces the aggregate crossover. Then we have two groups. Among those not exposed to tobacco, mortality is as given in Equations 1a-d, with f representing the mortality multiplier on having the cancer gene. Among those exposed to tobacco, these individual-level equations are the same but for a new mortality multiplier, which we can call t (for 'tobacco' or, equally, for the 'trait' that will be conditioned on). This is given in Eq. 6a-d:

$$\begin{aligned}
\mu_{w,r,t}(a) &= t\alpha e^{\beta a} \\
\mu_{w,f,t}(a) &= ft\alpha e^{\beta a} \\
\mu_{b,r,t}(a) &= bt\alpha e^{\beta a} \\
\mu_{b,f,t}(a) &= bft\alpha e^{\beta a}
\end{aligned} \tag{6}$$

What happens when an empirical analysis conditions on fetal tobacco exposure but does not observe the genetic predisposition? This may be what analysts are imagining when they condition on particular dimensions of heterogeneity.

One way to understand this situation is to see that, for each trait that affects individual-level heterogeneity, three processes are happening simultaneously, as summarized in the three rows of Table 1. The first and second processes encapsulate the way we think about *frailty* in the unidimensional context: frailty raises mortality at the individual level, and therefore it also is selected out of the cohort, particularly among blacks. Meanwhile, the first and third processes the way we think about *race* in the unidimensional context: being black raises mortality at the individual level, and therefore (among survivors) it also indicates robustness, since frail blacks are unlikely to survive. Yet in fact, *frailty* and *race* play completely symmetrical roles in the mortality equation; the difference in which of their dynamics catches our attention reflects merely that our interest is in explaining the black-white crossover, which is to say, that when we analyze the crossover, we condition on race but are unable to condition on frailty because it is unobserved.

For the dimension of heterogeneity that we *do* condition on, then, all three processes come into play. Fetal tobacco exposure, by assumption, raises mortality in individuals at all ages. It is

therefore selected out of the cohort, and disproportionately so selected among blacks. Fetal tobacco exposure is therefore also an indicator, among survivors—and especially among black survivors—of *lacking* the cancer gene. The logic is the same as the logic of the crossover itself: if one is unlikely to survive to an old age as while holding a genetic predisposition to cancer and also facing the disadvantages associated with being black in the United States, then one is even less likely to survive those circumstances if one *also* was exposed to tobacco in utero. Among those who were exposed to tobacco, mortality selection against the cancer gene is profound. This logic is indeed the root of the crossover: selective survival—the world conditioning on mortality—creates an association between race and frailty, even if none existed at birth. Our point is simply that the same process occurs, not only between frailty and race, but between the distinct dimensions of frailty as well.[16]

Process: T group compared to non-T group	**Effect on aggregate mortality**	**For which race?**
T group has higher mortality in individuals conditional on frailty	Raises aggregate mortality	Blacks and whites equally
T group is selected out of the cohort	Lowers aggregate mortality	Conditional on frailty, stronger selection against T for blacks than whites; aggregated over frailty, either race can be more selected against T
T group has intensified selection against frailty	Lowers aggregate mortality	Either race can have a larger frailty difference between T and non-T

Table 1. How T (e.g., fetal tobacco exposure), which raises mortality at all ages in individuals, affects aggregate mortality.

[16] In the language of causal inference theory, the association occurs because mortality is a *collider* for its various risk factors, which selection then conditions on. See Pearl (2009) for an elaboration of the concept of a collider.

Conditioning mortality on fetal tobacco exposure and comparing the age at crossover that results, either between the exposed and non-exposed (i.e., T and non-T) groups, or between one of those groups and the aggregate, is informative only if the mortality selection model makes a clear prediction about which of these populations should reach a crossover first. Unfortunately, it does not. We will first consider the T compared to the non-T group, and then compare both of those subpopulations to the aggregate.

Crossover order among subpopulations: T vs. ~T

In each of our two subpopulations, the higher-mortality T (e.g., those fetally exposed to tobacco) and the lower-mortality non-T (those not so exposed), selection against the unobserved cancer gene, at different rates for blacks and whites, might produce a black-white mortality crossover. If the higher mortality of the T group licenses a definitive prediction that that group will reach a crossover at either *older* or *younger* ages than the non-T group, then that prediction allows the mortality selection model to be falsified by empirical data, and therefore allows it to be tested.

Unfortunately, we will show that neither group always crosses before the other, suggesting that, if testable predictions about the order of reaching crossover are to be found, they must be based in a more specific set of assumptions about the workings of a selection model (e.g., specifying a precise range of parameter values, including for unobserved parameters) than have been put forward so far in the literature drawing on multiple dimensions of heterogeneity. We justify this conclusion first analytically and then with simulations.

Formal model

Equation 5 showed that every closed cohort (as defined by the respective mortality multipliers on being black and being frail, *b* and *f*) identifies a critical value which, in order for there to be a crossover, must be exceeded by the difference between the proportion frail among whites and the proportion frail among blacks times the black mortality multiplier. The age at crossover thus depends on the age at which this difference first exceeds this critical value.[17] This general fact about the age at crossover with a single dimension of heterogeneity applies to each of the T and non-T groups; the difference between them (as shown in Equations 6a-d, which simply modify Equations 1a-d with the random variable T) is simply a difference in intercept, reflected in the mortality multiplier *t*.

Therefore, one way to identify a consistent prediction about which group will first reach a crossover would be to show that the difference between, on the one hand, the proportion frail among whites and, on the other hand, the proportion frail among blacks times the black mortality multiplier is always larger among the T compared to the non-T, or vice-versa.[18] Thus, consider the sign of the expression in Equation 7:

$$\left(\pi_{wt}(a) - b\pi_{bt}(a)\right) - \left(\pi_{wn}(a) - b\pi_{bn}(a)\right) \tag{7}$$

[17] This equation also implies that cohorts will uncross when this difference falls below the critical value (as frail whites become sufficiently close to extinction that the proportion frail among whites reduces the left-hand side of Equation 5). That such uncrossings have not been empirically observed, to our knowledge, could be considered a refutation of this selection model, or it could be taken to reflect the limited data available at very old ages. (Even studies that include very old ages usually measure mortality with parametric assumptions over age that may or may not allow mortality to appear uncrossed after it has crossed. For example, studies that estimate the crossover with only a race interaction with a log-linear main effect of age will not allow uncrossings.)

[18] Such a consistent sign in the inequality between the two differences would be a sufficient, but not a necessary, condition for one group to always cross before the other: what would be strictly necessary would be that this inequality hold only in the ranges of values in which crossovers occur.

Being black, being frail, and being T each increase the intercept of mortality relative to being white, robust, and non-T, respectively; we can call that intercept α*, which equals α multiplied by the relevant mortality multipliers, *b*, *t*, and *f*. The problem for drawing conclusions about the sign of the expression in Equation 7 is that increasing α* decreases π(a) non-linearly, as shown in Equations 8 and 9. Equation 8 gives the proportion frail in each subpopulation, and is simply a different form of Equation 3b; we give it because it is the form that employs the simulation parameters, described below:

(8)

$$\pi_{g,race}(a) = \frac{\pi_g(0)\cdot\left[S(a)\right]^{f\cdot TB}}{\pi_g(0)\cdot\left[S(a)\right]^{f\cdot TB} + \left(1-\pi_g(0)\right)\cdot\left[S(a)\right]^{TB}}$$

where the g term denotes the proportion frail for either the T or the non-T group and T and B are set either to 1 (for the non-T and whites, respectively) or to *t* and *b*, respectively, for the T group and blacks.

Equation 9 gives the rate of change in the proportion frail:

(9)

$$\begin{aligned}\pi'(a) &= -\pi(a)\left(1-\pi(a)\right)\left(\mu_f(a)-\mu_r(a)\right)\\ &= \frac{-p\left(f-1\right)\alpha^* e^{\beta a}\omega_k}{\left(p+\omega_k\right)^2}\end{aligned}$$

The first line of Equation 9 was given in Wrigley-Field (2014), and is a special case of an equation given in Vaupel and Zhang (2010). The second line, which is new here, relates that form to the underlying mortality parameters. The *p* in the numerator, divided by the term that is squared in the denominator, is π(a). The exponential term on the right of the numerator (which

equals the ratio of the survivorship of the robust to the survivorship of the frail, as shown in Equation 3b), divided by the term that is squared in the denominator, is 1-π(a), or the proportion robust among the survivors to age *a*. And the term $(f-1)\alpha^* e^{\beta a}$ is the difference between frail and robust mortality.

Equation 9 allows us to rewrite Equation 7 into a form that may yield more intuition for its sign, as in Equation 10, using ω to represent the ratio of the survivorship of the robust to the frail among non-T whites, $\omega = \frac{S_r}{S_f}$, and μ(a) to represent the mortality of robust non-T whites:

$$\left(\pi_{wt}(a) - b\pi_{bt}(a)\right) - \left(\pi_{wn}(a) - b\pi_{bn}(a)\right) = p\mu(a)\left[\frac{b^2\omega^b}{\left(p+\omega^b\right)^2} + \frac{t\omega^t}{\left(p+\omega^t\right)^2} - \frac{\omega}{\left(p+\omega\right)^2} - \frac{b^2 t\omega^{bt}}{\left(p+\omega^{bt}\right)^2}\right] \quad (10)$$

The key implication of this equation is that there is no obvious generalization about the order of the crossover (i.e., the sign of the expression in Equation 7/Equation 10). Equation 10 suggests that whether the non-T or T is closer to the crossover critical value depends on whether the rate of the loss of frailty in the group that is both black and T is so great that it outweighs the degree of loss among the black non-T and the white T. That is: given the assumption of equal frailty in all groups at baseline, at subsequent ages, the level of frailty is lower among T whites than non-T whites, and non-T blacks than non-T whites; it is lowest of all among T blacks. The order of the crossover is therefore determined by which sum is larger: the frailty of the frailest (non-T whites) and the least frail (T blacks) group (with the latter weighted by *b*), or the frailty of the two in between (T whites and non-T blacks, with the latter weighted by *b*).

The form of the inequality suggests that its sign will vary with the levels of frailty involved. The intuition is that the additional selection produced by a group's being T or black (or both) will cause the biggest reduction in the group's proportion frail when the level of the proportion frail is close to .5, when the rate of selection is particularly high (Wrigley-Field 2014). Intuitively: if the selection among the white T group has been sufficiently large, so that the proportion frail among white survivors is very small (close to zero), then even though the proportion frail among black survivors will be even smaller, the difference between the two is limited. Conversely, if the selection among the black non-T group has been sufficiently small that their proportion frail is still relatively high (close to 1, which requires that it was very high at baseline), then, even though the proportion will be even higher among white survivors, the difference between the two is limited.[19]

In short, the equations (as we have analyzed them) do not seem to us provide any *a priori* basis for concluding that either the T group or the non-T group will always reach a crossover before the other, and our intuitive interpretation of Equation 10 suggests (but does not prove) that this may vary. The simulations presented next do prove that this is so: in either the T or the non-T group, black and white hazards can cross without the hazards in the other group having crossed previously.

[19] These observations are compatible with the T always reaching crossover before the non-T (though not the reverse), since the T groups must pass through larger values of the surviving proportion frail to reach the smaller values just discussed. But it does not follow that at any point at those larger values, the difference between the white and black proportion frail will be sufficient to produce a crossover among the T. Such complexities in reasoning intuitively and analytically motivate the recourse to simulations.

Simulation procedure

The simulations analytically calculate the existence or absence of a crossover in 34,992 distinct cohorts defined by six parameters: the mortality multipliers on being black, T, and frail; the baseline proportion frail among the T and non-T; and the baseline proportion T. All three proportions are assumed to be equal at baseline for blacks and for whites, and are modeled in the range .55 to .95 in units of .05 (a range chosen so that crossovers might be viewed when the frail are either a majority or—as their share among survivors declines—a minority in the various subpopulations). To maximize the probability of there being crossovers whose timing we can compare, we use small values of the mortality multiplier on being black—1.05, 1.5, and 2—and large values of the mortality multipliers on being frail and T—all even numbers between 2 and 8. Our purpose is not to yield a substantively realistic mortality model (which after all is not likely with only two dimensions of heterogeneity, both fixed at birth), but to illustrate the substantial complexities that arise when heterogeneity is any more than unidimensional.

Each of the 34,992 cohorts is examined at 190 different age units, defined by the baseline survivorship, as it declines from .995 to .05 in units of .005, yielding 6,648,480 total observations that might or might not be crossed at any group. That is, we look for crossovers in each cohort as the proportion of its non-T robust whites that survives goes from being nearly everyone to being only 5% of the original subcohort (at which point the other, higher-mortality groups are essentially extinct), in units of half of one percent of survivorship.

Since survivorship declines monotonically, but not linearly, with age, the baseline survivorship gives an ordinal, but not a cardinal, measure of age within each cohort: it tells us which crossovers happen before or after which other ones. By thus ordering by survivorship

instead of by age per se, we avoid the need to choose values for the baseline mortality intercept and the mortality log-slope for each group (since each survivorship value could be reached by infinitely many combinations of intercept, slope, and age value). Instead of simulating full mortality models, we can exploit that fact that all cohorts and subcohorts, as they age, move monotonically from a survivorship of 1 to a survivorship of zero, using the formulas for the crossover point in terms of fixed parameters and the race-specific proportions frail (given in Equations 3 and 5 for the unidimensional case and Equation 8 and Equation 12, below, for the aggregate cohort). This simplifies the simulations in computation and interpretation without losing any information relevant to our question: which group crosses first?

Simulation results

We find that the T group can reach a black-white crossover at a younger age than the non-T group crosses, and that blacks and whites in either group can cross without the other ever crossing. We do not find that the non-T group crosses at a younger age and the T group at an older age, in these simulations.

In 12,636 simulated cohorts, or 36% of the total, the T group crosses either before the non-T group crosses, or with the non-T group never crossing. This occurs across the full range of simulated values of the frailty multipliers on being T and being black, and at values of the frailty multiplier of at least 4. The critical value in these cohorts as defined by Equation 5—the ratio of the extent to which black mortality exceeds white and the extent to white frail mortality exceeds robust—ranges between .007 and one-third. That critical value is the value that, for a crossover to occur, must be exceeded by the difference between, on the one hand, the proportion frail

among whites, and on the other, the proportion frail among blacks times the black mortality multiplier. In 60% of these cohorts, the non-T do eventually cross (and in 80%, the aggregate group, discussed below, does as well).

In 5,112 simulated cohorts, or 15% of the total, the non-T cross without the T having crossed beforehand.[20] This occurs in the same range of values of the mortality multipliers the critical value as the reverse, early cross among the T, occurred. In none of these simulated cohorts does the T group eventually cross (in 70% of them, the aggregate cohort does).

Crossover order between subpopulations and the full population

The mortality models for the T and non-T groups are each just instances of the same unidimensional frailty case that the mortality selection literature has carefully analyzed. But the dynamics of the multidimensional case are wholly different. We begin by asking which group will reach crossover first: the T and non-T subpopulations, those who were and were not fetally exposed to tobacco, or the whole cohort, aggregating over fetal exposure.

Formal Model

In contrast to the two-group case, with two dimensions of heterogeneity, aggregate mortality reflects that there are four groups, defined by having either, both, or neither of our two traits, fetal tobacco exposure and the cancer gene. Let $\pi_n(a)$ be the proportion frail among the non-T at some age, and $\pi_t(a)$ be the proportion frail among the T. Let $z(a)$ be the proportion T

[20] In the remaining 49% of cohorts, neither the T nor the non-T groups cross.

among the total race-specific cohort. Then the mortality of this cohort at any age is given by Equation 11:

$$\bar{\mu}(a) = B \cdot \mu(a) \cdot \begin{pmatrix} ft \cdot z(a)\pi_t(a) + t \cdot z(a)\left(1-\pi_t(a)\right) \\ +f \cdot \left(1-z(a)\right)\pi_n(a) + \left(1-z(a)\right)\left(1-\pi_n(a)\right) \end{pmatrix} \tag{11}$$

where B=b for blacks and B=1 for whites, and μ(a) gives the mortality of robust non-T whites at that age (a mortality level that is scaled up for other groups by the mortality multipliers).

The proportion T, z(a), is given in Equation 12 as:

$$z_{race}(a) = \frac{z(0)\cdot\pi_t(0)\cdot\left[S(a)\right]^{ft\cdot B} + z(0)\cdot\left(1-\pi_t(0)\right)\cdot\left[S(a)\right]^{t\cdot B}}{\begin{pmatrix} z(0)\cdot\pi_t(0)\cdot\left[S(a)\right]^{ft\cdot B} + z(0)\cdot\left(1-\pi_t(0)\right)\cdot\left[S(a)\right]^{t\cdot B} + \\ \left(1-z(0)\right)\cdot\pi_n(0)\cdot\left[S(a)\right]^{f\cdot B} + \left(1-z(0)\right)\cdot\left(1-\pi_n(0)\right)\cdot\left[S(a)\right]^{B} \end{pmatrix}} \tag{12}$$

With these expressions for aggregate mortality and the proportion of each race that is T, and the expressions given in Equation 8 for the proportion of each race and T/non-T group that is frail, we can fully characterize the crossover. As in the unidimensional case, we can provide the conditions for the crossover in terms of the mortality multipliers and the population compositions (which themselves are expressed in terms of the baseline compositions and the lowest-mortality group's survivorship at age *a*, precluding the need to model mortality directly in the simulations). Equation 13 gives the condition for the crossover in aggregate mortality:

$$\begin{aligned} \bar{\mu}_b - \bar{\mu}_w < 0 \Leftrightarrow \\ b \cdot z_b(a)\Big[\big(f-1\big)\big(t \cdot \pi_{tb}(a) - \pi_{nb}(a)\big) + t - 1\Big] + b \cdot \Big[\big(f-1\big)\pi_{nb}(a) + 1\Big] \\ < z_w(a)\Big[\big(f-1\big)\big(t \cdot \pi_{tw}(a) - \pi_{nw}(a)\big) + t - 1\Big] + \Big[\big(f-1\big)\pi_{nw}(a) + 1\Big] \end{aligned} \tag{13}$$

One rough way to get some intuition for this form is to ask when each pair of terms (the first term or the second term on each side) will be smaller on the left-hand side.[21] For the rightmost pair of terms of Equation 13, rearrangement reveals that the rightmost term on the top (black mortality) line to fall below the rightmost term on the bottom (white mortality) line exactly when there is a black-white crossover among the non-T group, shown in Equation 14:

$$\begin{aligned} b \cdot \Big[\big(f-1\big)\pi_{nb}(a) + 1\Big] < \Big[\big(f-1\big)\pi_{nw}(a) + 1\Big] \Leftrightarrow \\ \frac{b-1}{f-1} < \pi_{nw}(a) - b \cdot \pi_{nb}(a) \end{aligned} \tag{14}$$

This, of course, is just Equation 5, applied to the non-T group. In other words, when black mortality falls below white mortality among the non-T, driven by a large disparity in the white and black proportions of unobserved frailty, then this will tend to contribute to a racial crossover in the whole cohort as well.

Comparing the leftmost terms in Equation 13 is a bit more complex than the rightmost terms just compared. The dynamics of these leftmost terms are easier to understand, in relation to what we have seen already, when slightly rearranged, as in Equation 15:

[21] If the pairs had the same direction of inequality, then of course that would be the inequality of the whole term, and would tell us whether there is a crossover; but if they have different signs, then the magnitudes of the relevant differences determine the direction of the overall inequality, that is, whether there is a crossover. Looking at the terms in two separate inequalities, to see when each pair of terms contributes to a crossover, is just a way of understanding what the full inequality is saying.

$$
\begin{aligned}
& b \cdot z_b(a)\left[(f-1)\left(t \cdot \pi_{tb}(a)-\pi_{nb}(a)\right)+t-1\right] \\
& <z_w(a)\left[(f-1)\left(t \cdot \pi_{tw}(a)-\pi_{nw}(a)\right)+t-1\right] \\
& \Leftrightarrow \\
& b \cdot\left[\left(t \cdot \pi_{tb}(a)-\pi_{nb}(a)\right)+\frac{t-1}{f-1}\right] \\
& <\frac{z_w(a)}{z_b(a)} \cdot\left[\left(t \cdot \pi_{tw}(a)-\pi_{nw}(a)\right)+\frac{t-1}{f-1}\right]
\end{aligned}
\tag{15}
$$

The first step to understanding Equation 15 is to see that the bracketed set of terms on each side is negative precisely when the T and non-T groups, of the respective race, have had their own race crossover—that is, whether the average mortality of the T has fallen bellow the average mortality of the non-T due to greater selection against frailty among the T. This is shown in Equation 16, which has the analogous form to Equation 5, but with *t* in place of *b*:

$$
\begin{aligned}
& t \cdot \pi_{t,race}(a)-\pi_{n,race}(a)+\frac{t-1}{f-1}<0 \Leftrightarrow \\
& \pi_{n,race}(a)-t \cdot \pi_{t,race}(a)>\frac{t-1}{f-1}
\end{aligned}
\tag{16}
$$

We can therefore conceptualize the full set of terms of the inequality represented in Equation 15 with reference to a two-by-two table, reflecting whether the T and non-T groups have had a mortality crossover among blacks, whites, both, or neither. This is shown in Table 2.

		Whites	
	Does T mortality fall below non-T mortality among...	NO	YES
Blacks	NO	LHS > 0, RHS > 0 This set of terms can either contribute to or mitigate the tendency to aggregate crossover. Here, $b < z_w/z_b$ contributes to the crossover.	LHS > 0, RHS < 0 This set of terms always mitigates the tendency to aggregate crossover
	YES	LHS < 0, RHS > 0 This set of terms always contributes to the crossover	LHS < 0, RHS < 0 This set of terms can either contribute to or mitigate the tendency to aggregate crossover. Here, $b > z_w/z_b$ contributes to the crossover.

Table 2. Each cell gives the sign of the left-hand side (LHS) and right-hand side (RHS) of the inequality in Equation 13, and says whether the set of terms in Equation 13 contribute to (or mitigate against) an aggregate crossover.

As reflected in Table 2, two factors determine whether the inequality in Equation 13 holds, contributing to a crossover in the aggregate. First is the question of whether the hazards of the T and non-T groups have crossed for each race, which determines the sign of the bracketed terms. Second, if the right and left sides of the inequality have the same sign, as in the upper left and lower right cells of Table 2, is the question of which one is larger. This, in turn, is determined partly by the relative magnitudes of the mortality multiplier on being black, *b*, relative to the ratio of the proportion of the surviving cohort that is T in whites relative to blacks. When neither race has had a crossover between the T and non-T groups, then the full set of terms is more likely to contribute to a crossover when the black mortality multiplier is smaller than the white/black

ratio of the proportion T. When both races have had a T/non-T crossover, then conversely, the full set of terms is more likely to contribute to a crossover when the black mortality multiplier is *larger* than the white/black ratio of the proportion T—reflecting that, if the T and non-T have crossed hazards, having a greater proportion of T members will now tend to lower aggregate mortality, whereas otherwise it would raise aggregate mortality. When such a T/non-T crossover has occurred only among blacks, this set of terms will always contribute to an aggregate white-black crossover; when such a T/non-T crossover has occurred only among whites, this set of terms will always work against an aggregate white-black crossover.

Taken together, these equations are suggestive—but only suggestive—that an aggregate mortality crossover might occur either before or after the underlying T and non-T subpopulations cross. To be certain, we return to simulations.

Simulation Results

Analyzing the same 34,992 simulated cohorts (each observed at 190 ages) as before, we focus on two possible outcomes: can the aggregate black and white subcohorts reach a crossover earlier than both of the underlying T and non-T groups? And can the aggregate reach a crossover later than both of the T and non-T groups?

The answer to both questions is yes. The aggregate crosses first in 10,752 simulated cohorts, or 31% of the total. These are circumstances in which an analyst, conditioning on T, would find the crossover moving to an *older* age than it did in the aggregate. These outcomes occur in simulated cohorts across the range of values of the mortality multipliers on being frail and being T, and at crossover critical values ((b-1)/(f-1)) ranging from close to zero to 1.

Eventually, 26% of these cohorts will have a crossover in their non-T group, and 40% in their T group.

In these cohorts, by definition, at the time of the aggregate crossover, neither the T nor the non-T groups has reached a crossover. What, then, drives the crossover in the aggregate? Because of the complexity of the condition for an aggregate crossover, there is no single answer to this question, but some remarks help to illustrate how the crossover may arise from the combination of the racially differential selection against T with the differential selection against frailty within the smaller groups defined by T/non-T status and race.

In terms of the possibilities identified in Table 2, at the time that the aggregate group reaches its crossover, we are in the upper left quadrant of the table: only 34 of these cohorts (3%) ever show such a crossover in either race (all 34 eventually have a T/non-T crossover for both blacks and whites), and none of them are a T/non-T crossed state for either race at the time that the aggregate cohort crosses. In the absence of a race crossover among the non-T and any crossovers between the T and non-T, then, following from Equation 13, one relevant factors in producing an aggregate crossover are the extent to which the white-black ratio of the proportion T exceeds the black mortality multiplier ($b < \frac{z_w(a)}{z_b(a)}$). In 64% of these cohorts, the white-black ratio of the proportion T is indeed larger than the black mortality multiplier at the time of crossover. Another relevant factor is the extent to which whites are frailer than blacks, aggregated over T.[22] At the time that aggregate mortality crosses, in all but 1% of these cohorts,

[22] Since T and frailty play symmetric roles in the mortality model, Equation 13 could be rewritten in terms of the white/black ratio of the proportion frail instead of the ratio of the proportion T. But the interpretation of this ratio would be equally complex because the frail group is selected against being T, just like the T group is selected against being frail, and this selection also happens more sharply for blacks than for whites.

the white-black imbalance in the proportion frail is larger in the aggregate than it is in at least one of the T or non-T groups (but it is only larger than the imbalance in *both* the T and non-T groups in 7.5% of cohorts). In general, then, these are the cohorts in which conditioning on T tends to also make the black and white groups reasonably more similar in terms of unobserved frailty as well.

But the opposite crossover order is possible as well. In 3,413 cohorts (10% of the total), both the T and non-T groups cross before the aggregate does. In contrast to the cohorts just considered, these are circumstances in which an analyst, conditioning on T, would find the crossover moving to a *younger* age. These outcomes occur across the full range of values of the mortality multiplier on being T, and at values of the frailty multiplier of at least 4. The critical crossover value in these cohorts ranges from close to zero to 1/3. Eventually, aggregate mortality will also cross in 66% of these cohorts.

What is preventing the aggregate crossover at the time that the second group (always the non-T, since a non-T cross is never followed by a T cross) crosses? Again, we can provide some suggestive remarks. In terms of the framework provided by Table 2, in 68% of these cohorts, at the time of the non-T cross, neither race has had a T/non-T crossover; we are in the upper left quadrant of Table 2. But 16% of the cohorts have had a T/non-T crossover in both racial groups, meaning that a larger proportion T now tends to lower mortality, not raise it. In all of these cases, however, the white/black ratio of the proportion T is below 1: blacks have a larger proportion T than whites do.[23] In terms of the full conditions for the aggregate crossover given in Equation 13, then, most of them are met: the non-T have a race crossover; the relative sizes of the T groups

[23] It may be counterintuitive that this is possible. The next section shows why it is.

contribute to a crossover. It must be the case that the magnitudes of the bracketed terms in Equation 15 have a larger absolute value for blacks than whites—and indeed, this is the case.

Perhaps more tellingly, in the remaining 16% of these cohorts, whites but not blacks have a T/non-T crossover when the non-T cross, putting us in the upper right quadrant of the table—the quadrant in which the set of terms in Equation 14 always tends to mitigate the aggregate crossover, perhaps accounting for why it does not occur. Similarly, in none of these cohorts have only blacks had a T/non-T crossover: at the time that the non-T cross, we are never in the lower left quadrant of Table 2, in which the terms in Equation 15 always tend to produce a crossover. This helps to explain the absence of an aggregate crossover.

In all of the cohorts in which both subgroups cross before the aggregate does, at the time that the non-T cross, the white-black frailty difference is larger in the non-T (though not the T) than it is in the aggregate. Conditioning on T, then, *exacerbates* the frailty imbalance, relative to the non-T group—presumably the opposite of what most analysts conditioning on a variable might wish for. In the next section, we more explicitly consider the percent frail in the aggregate black and white subcohorts.

Frailty Reversals, or: Why multidimensional frailty is fundamentally different

Our analysis of the age at crossover in the race subcohorts aggregated over T has made no direct reference to the proportion frail in those aggregated groups of whites and blacks. Here we explicitly consider the aggregate proportion frail among whites and blacks, with a surprising result.

In the unidimensional heterogeneity model, a simple generalization about frailty drives the crossover: at ages above baseline, blacks always have a lower proportion frail than whites. For understanding the crossover, this is *the* key fact about mortality selection.

But it turns out that when there is more than one dimension of heterogeneity, this basic generalization—whites have more surviving frail members than blacks—need not hold. When there are two dimensions of heterogeneity, mortality selection can lead blacks to have either *lower* or *higher* composition of either one individually. In the multidimensional selection model, mortality selection can make blacks frailer than whites, if "frailty" means (as it traditionally does) that part of heterogeneity which goes unmeasured.[24] In what follows, we exploit the generalization about the unidimensional model, which applies to the T and the non-T groups, to show how it may fail in the aggregation of the two.

Formal Model

The proportion frail among the *aggregate* population of whites or blacks, which we denote θ, is given in Equation 17 in terms of the proportion of the group that is T, $z(a)$, and the proportion frail among the T and the non-T:

[24] Note, however, that this result holds only in the population aggregated over even the part of the heterogeneity that is measured. Put differently, the result is that any individual dimension of heterogeneity, rather than the entirety of the heterogeneity, might be made more prevalent among whites than blacks, solely because of mortality selection.

(17)

$$\theta_{race}(a) = \pi_{t,race}(a)\cdot z_{race}(a) + \pi_{n,race}(a)\cdot\left(1 - z_{race}(a)\right)$$
$$= \frac{\pi_t(0)\cdot z(0)\cdot S^{ftB} + \pi_n(0)\cdot\left(1-z(0)\right)\cdot S^{fB}}{\left(\begin{array}{l}\pi_t(0)\cdot z(0)\cdot S^{ftB} + \pi_n(0)\cdot\left(1-z(0)\right)\cdot S^{fB} \\ +\left(1-\pi_t(0)\right)\cdot z(0)\cdot S^{tB} + \left(1-\pi_n(0)\right)\cdot\left(1-z(0)\right)\cdot S^{B}\end{array}\right)}$$

The first line is far simpler; the second line serves as a reminder that the variation in survivorships across the difference groups drives the resulting proportion, and that that variation in survivorships is the only thing that varies by race (reflected in the B term, where B=1 for whites and B=*b* for blacks).

Can whites have a *lower* proportion frail than blacks? Equation 18 gives the condition for this 'pathological case' to be true:

(18)

$$\theta_w(a) < \theta_b(a) \Leftrightarrow$$
$$\pi_{t,w}(a)\cdot z_w(a) + \pi_{n,w}(a)\cdot\left(1-z_w(a)\right) < \pi_{t,b}(a)\cdot z_b(a) + \pi_{n,n}(a)\cdot\left(1-z_b(a)\right) \Leftrightarrow$$
$$\pi_{n,w}(a) - z_w(a)\left(\pi_{n,w}(a) - \pi_{t,w}(a)\right) < \pi_{n,b}(a) - z_b(a)\left(\pi_{n,b}(a) - \pi_{t,b}(a)\right)$$

The second form of this condition, shown in the third line, is the most revealing. First, to the extent that the white non-T group has a larger proportion frail than the black non-T group (and it always will if the two groups had the same proportion frail at baseline)[25], this will tend to also make whites more frail than blacks in the aggregate, tending to prevent the pathological

[25] This is a more restrictive assumption than we made in the general equations for aggregate mortality, which stipulated only that all baseline proportions were equal across race, not across T status.

case. In order for blacks to nevertheless be more frail in the aggregate, then, the difference in the proportion frail between the non-T and the T, times the proportion of the cohort that is T, must be substantially larger for whites than for blacks.

Second, that difference in the proportion frail between the non-T and the T is always positive: for each race, the non-T are always less frail than the T. The result of the positive sign is that, since the second term on each side of the inequality is subtracted from the first, the proportion of each race that is T will tend to reduce the race's aggregate frailty, since the T are less frail than the non-T due to mortality selection. This fact carries the key intuition to how mortality selection might make blacks more frail than whites: if mortality selection makes whites more T than blacks, and the T less frail than the non-T, then it might sometimes thereby make whites less frail than blacks. If this were the case, then the z(a) term would be larger for whites than for blacks,[26] and the difference between the non-T and the T proportion frail might be larger for either whites or blacks,[27] but would have to be large enough for at least one of them to overcome the fact that the white non-T group is more frail than the black non-T group.

This reasoning suggests that it *might* be the case that blacks are sometimes made *frailer* than whites by mortality selection, but it does not prove it. Can it be simultaneously the case that the proportion T is sufficiently larger among whites than blacks, the difference between the non-T and the T proportion frail sufficiently large among at least one of whites or blacks, and the

[26] We guess (but have not shown) that whites will generally, but not always, be more T than blacks—by analogous reason to the argument made here about frailty—but that blacks cannot simultaneously be more frail and more T under this model. One or the other, but not both.

[27] We believe this difference can be larger either among whites or among blacks, depending on which has frailty proportions closer to .5, by reasoning analogous to that given above, when we compared the age at crossover in the T compared to the non-T. This part of the inequality in Equation 18 is an analogue of the inequality on the left-hand side of Equation 10, except for the absence of the mortality multiplier; what has changed is just whether we are comparing the difference between blacks and whites in T vs. non-T, or the difference between T and non-T across race.

white-black difference in frailty among the non-T sufficiently small, to create this pathological case? We use simulations to find out.

Simulation Results

A "frailty reversal," in which blacks have greater aggregate mortality than whites for some period of time, occurs in 14% of simulated cohorts. These cohorts occur across the range of mortality multipliers such that the multiplier on being T is larger than the multiplier on being frail (so in practice, given the simulated ranges, $f \leq 6 \,\&\, t \geq 4$). Within this sample of cohorts with at least one frailty reversal, a mean of 21% of the age observations are ones in which frailty is reversed. The number of age observations with a frailty reversal per cohort ranges from 1 to 108 (57% of the simulated age observations).[28]

The conditions allowing this pathological case are simplest to understand, as the discussion of Equation 18 indicated, when the T and non-T groups had the same baseline frailty, since that allows us to be sure that, for each race, the non-T group has more frail members than the T group. There are 423 cohorts with this equal baseline frailty constraint that experience a frailty reversal, representing only 1% of the total simulated cohorts, but 11% of the cohorts with equal baseline frailty for the T and non-T.

What is happening when frailty is reversed, with higher frailty among blacks, in these cohorts with equal baseline frailty between the T and non-T? As expected, when frailty is reversed, whites always have a greater proportion T than blacks do, as shown in Figure 1. Also as expected, the frailty difference between the non-T and the T, which is always positive (for a

[28] These numbers are only meant to provide rough descriptives, since the age measurement is only ordinal, not cardinal. This means that, were we to translate S into a linearly increasing age unit (e.g., years), the proportions would be different.

specific race), may be either larger or smaller for whites, even in these pathological cases. When it is smaller, the extent of the larger proportion T among whites is enough to counteract every other term in Equation 18 in generating smaller frailty for whites. Figure 2 shows together these two white-black differences in the terms subtracted from each side of Equation 18.

In summary, the simulations confirm the general impression created by Equations 17-18 and show how mortality selection can make blacks *more* frail than whites, with respect to individual dimensions of heterogeneity.

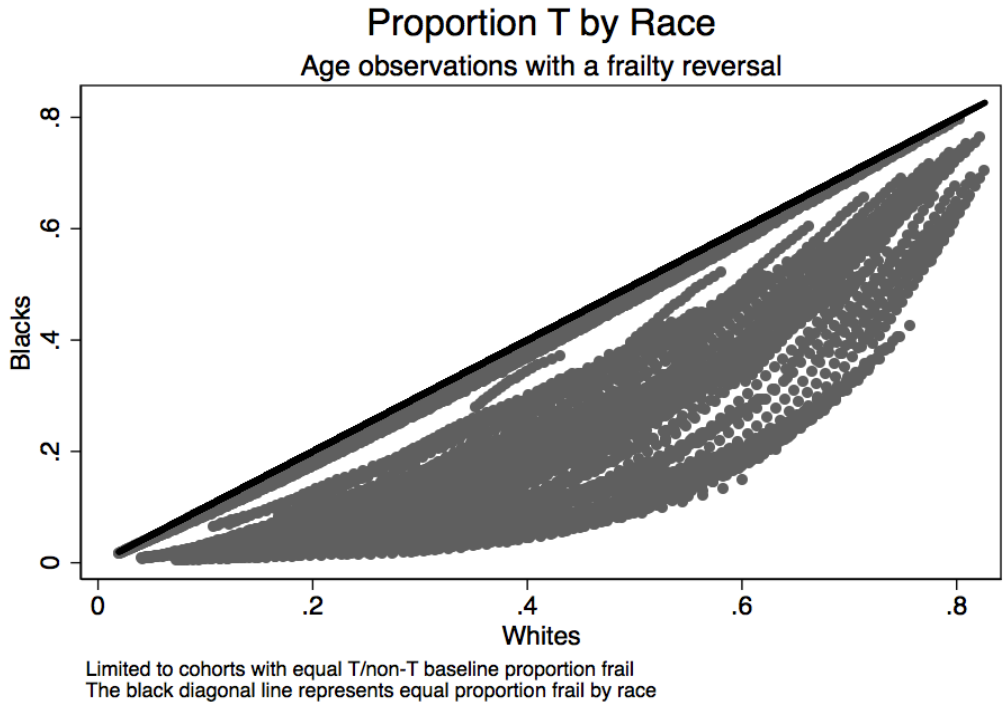

Figure 1. The proportion T for whites (x axis) and blacks (x axis) at observations with a frailty reversal and equal baseline frailty for all groups.

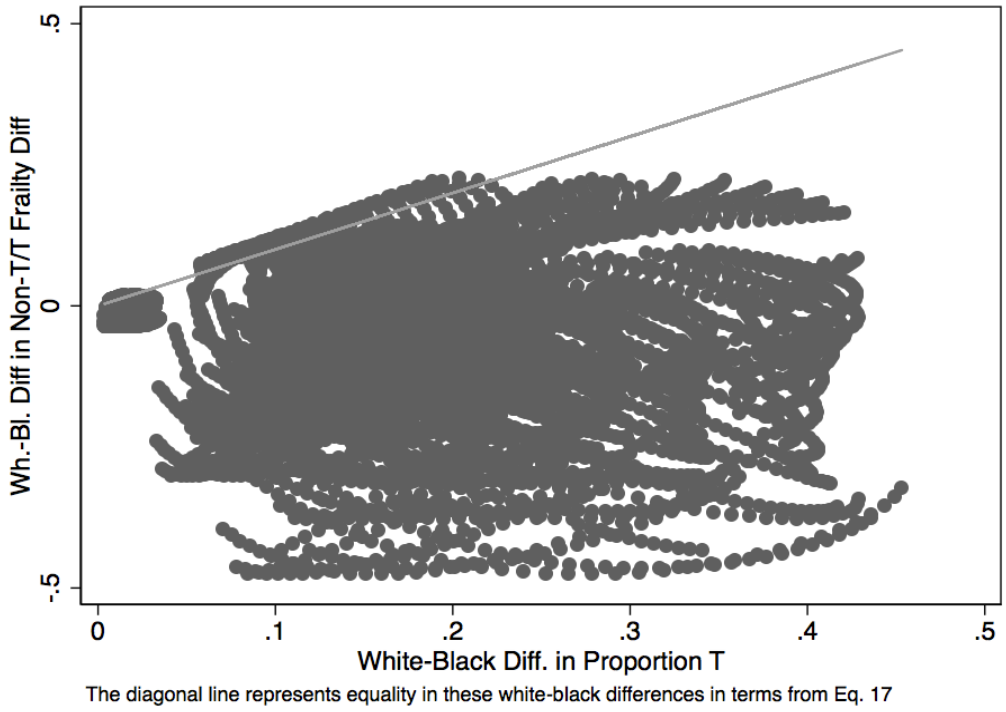

Figure 2. The white-black difference in the proportion T (x axis) and the white-black difference in the frailty disparity between the non-T and the T (y axis) at observations with a frailty reversal and equal baseline frailty for all groups.

Dynamic Heterogeneity

In connecting a mortality selection model like the one we have been discussing to the empirical work that conditions mortality on such dimensions of heterogeneity as socioeconomic status (e.g., Sautter et al. 2012, Yao and Robert 2011), an obvious question arises: If low socioeconomic status increases mortality so much, and if blacks are so generally selected due to their high mortality, why have all the low-SES blacks not been selected out of cohorts by old age? Indeed, how can poverty be *more* prevalent among old-age blacks than old-age whites?

The analysis we just provided, which showed how mortality selection can give blacks a *larger* share than whites of some specific dimension of heterogeneity, in principle constitutes one

answer. But a far simpler answer immediately presents itself: blacks are more likely to be poor at old age (in spite of mortality selection) because blacks are more likely to *become* poor at all ages, including old ages, leading the stock of impoverished blacks to be constantly replenished, even as it is diminished by mortality. In this sense, the dimensions of heterogeneity that have actually been conditioned on in the crossover literature are fundamentally unlike the fetal tobacco exposure of our stylized example, which is fixed by the time of birth.

That blacks do (we presume) become impoverished at higher rates than whites, even at relatively old ages, makes it unsurprising on its face that, as Sautter et al. (2012) and Yao and Robert (2011) find, conditioning on various dimensions of socioeconomic status tends to move the mortality crossover to younger ages: equalizing socioeconomic status, where otherwise blacks would face a disadvantage that raises their mortality, therefore reduces their mortality relative to whites, yielding a younger crossover.[29] No recourse to mortality selection at all is therefore needed to explain that the crossover occurs earlier, conditional on SES than in the aggregate.

On the other hand, once we acknowledge that low SES is not a fixed characteristic but one that is constantly being acquired, it can become problematic to speak of the low-SES group as intensely selected (which is the assumption underpinning the theoretical interpretations of these recent empirical analyses). In terms of our previous model, the movement of new people into the low-SES status at older ages—in the language of our model, people who become T—alters both the aggregate mortality of their racial group and the distribution of mortality across T and non-T

[29] This, of course, presumes that the low-SES group remains the higher-mortality group, though our previous analysis of multidimensional heterogeneity suggests that it need not, if the low-SES group is sufficiently selected by mortality. Indeed, Sautter et al. (2012) report an income crossover for men, albeit at older ages than the black-white crossover.

groups. First, this movement transforms people who are non-T into people who are T, which, conditional on frailty, raises mortality for those individuals, and thus the aggregate racial group. Second, these newly T have not been as heavily selected by mortality as the longstanding T, so, if their becoming T was independent of their frailty, they may be more frail on average than the longstanding T. This, in turn, will tend to reduce the frailty disparity between the T and the non-T, which may also alter crossover timing between the two groups.

Each of these effects serves to mitigate an effect of mortality selection, either in the aggregate racial group or in the T group. If blacks are especially disadvantaged in becoming T late in life, this may mitigate selection for blacks particularly, delaying the crossover.

There is a small literature on dynamic heterogeneity models in the context of mortality that adds precision to some of these observations. Mohtashemi and Levins (2002) present a dynamic unidimensional frailty model in which, translated into our language, both both blacks and whites can change from being robust to frail, and vice versa. They show how a crossover can occur either from elevated black mortality conditional on frailty, yet be offset by transitions from robustness to frailty among both blacks and whites, as we just discussed; or, alternatively, from blacks and whites having equal mortality conditional on frailty while blacks begin with more frail members but have a slower rate of becoming frail.[30]

A broader, if small, literature, not focused on the crossover, has engaged with status transitions either as an alternative to mortality selection (Rogers 1992), or a complement to it. A series of papers by Manton, Woodbury, and Stallard (Manton et al. 1994, 1995; Woodbury and Manton 1986) models mortality as a function of unobserved fixed frailty and observed dynamic

[30] This second model could be considered a dynamic, heterogeneous generalization of the model discussed in Liu and Witten (1995), which otherwise has no heterogeneity. Here, the movement between groups would provide the mechanisms for the racial patterns that Liu and Witten posit.

physiological states. Those physiological states are, in turn, modeled as functions of previous physiological states, independent of frailty, yielding a sophisticated model of progressive physical decline in the context of unobserved heterogeneity.

If future work on the crossover continues to condition on dimensions of heterogeneity that clearly are acquired during the life course while positing unobserved heterogeneity operating alongside, these dynamic models offer a foundation for new work on the crossover. But these dynamic models are surprisingly little discussed in the literature on the crossover. (For example, the Mohtashemi and Levins (2002) paper on the crossover, which also usefully discusses how to translate period crossovers into cohort crossovers—thereby providing a crucial link between most empirical datasets and mortality selection theory—is not cited in any of the other crossover papers we cite.) Yet they speak more directly than static models to the kinds of heterogeneity of greatest concern in recent demographic analyses.

Conclusion

Recent work on the black-white mortality crossover often conditions on some dimension of heterogeneity—such as religious participation (Dupre et al. 2006), income and education (Sautter et al. 2012), or neighborhood (Yao and Robert 2011)—and estimates the age at crossover. Seemingly regardless of what the analysts find about the order of the crossover in the resulting subpopulations and the aggregate population, they conclude that mortality selection explains the crossover, and that the dimension of heterogeneity they conditioned on is an important part of how that occurs.

Yet none of this recent work presents an explicit model of the mortality selection process that is to be evaluated. In this article, we formulated explicit mortality selection models under a variety of assumptions about heterogeneity and investigated what happens to the age at crossover when mortality is conditioned on one dimension of that heterogeneity. Our presumption in doing so was that, for the age at crossover to license some conclusion about mortality selection, a mortality selection model would need to make some prediction (which could be falsified or not) about the order of crossover between some groups: either between the subpopulations that result from conditioning (e.g., the poor and the non-poor), or between one of those subpopulations and the aggregate population.

If there is a single dimension of fixed heterogeneity ('frailty'), and selection against this single dimension fully accounts for the crossover, then conditioning on that dimension removes the crossover entirely by making black and white hazards log-parallel. Thus, in evaluating this model, if the black and white hazards become log-parallel after conditioning on heterogeneity, the mortality evidence would be consistent with the model; any other result would refute this model. But this model, although it is the basic of much theoretical work (e.g., Vaupel et al. 1979, Vaupel and Yashin 1985) that has served to make mortality selection vivid to demographers and develop intuition about how it works, is not (we presume) the model that anybody wants to test. The very exercise of conditioning on a particular, measurable dimension of heterogeneity seems predicated on the plausible assumption that the heterogeneity implicated in mortality selection theory has more than one dimension.

Therefore, the bulk of our paper explores the age at crossover under multidimensional heterogeneity. If there are two fixed dimensions of heterogeneity, and one is conditioned on

while the other is not (perhaps because it is unobserved in mortality data), then it turns out that either of the resulting subpopulations can cross while the other has not yet crossed, and further, both subpopulations can cross either before or after the aggregate population. In short, any order of crossover appears to be possible in the absence of more restrictive assumptions about the underlying values of the parameters, including latent parameters.

This suggests that any reported finding in this literature about the age at crossover is consistent with such a multidimensional fixed heterogeneity model. It also suggests that if every one of these papers had found the exact opposite order of crossover that they actually found, those results would be equally consistent with the multidimensional fixed heterogeneity model. The empirical tests of selection models used in this recent literature, then, are not really tests at all.

The results presented here broadly caution demographers against transporting insights derived from a unidimensional model into a multidimensional setting. For example, we show that, in terms of the unobserved dimension of heterogeneity, blacks may end up less selected, in total, than whites, precisely because they may be more selected along the observed dimension, and each dimension intensifies or mitigates the degree of selection on the other.

Finally, if heterogeneity is not just multidimensional, but also dynamic—if the heterogeneity that is conditioned on can be gained or lost over the life course, as in fluctuating characteristics such as income—then the selection dynamics may be undermined by this movement between groups. Without an explicit model that commits an analysis to any real assumptions about this dynamic process, there is even less basis for the kinds of clear predictions

about the age at crossover that would motivate the empirical exercise of evaluating that age under various statistical conditioning.

There is no question that the world in which blacks and whites live is riven by dynamic, multidimensional heterogeneity. It is a virtue of the recent crossover literature to have given renewed emphasis to relating the crossover to some of these myriad, intersectional inequalities within and between black and white cohorts. But it is a significant weakness of this recent literature to have at the same time abandoned the demographic tradition of explicit modeling of population processes. It is perhaps the core lesson of demographic theory that the reflection of individual-level dynamics in the population aggregates of cohorts or periods might take surprising shape. Just as the mortality of homogeneous cohorts is a poor guide to the mortality of heterogeneous cohorts divided along a single dimension (Vaupel and Yashin 1985), so the mortality of a cohort with two subpopulations is a poor guide to the mortality of a cohort with two cross-cutting heterogeneities. The irony of the crossover literature is to have become far more specific about what kinds of inequalities the *heterogeneity* of mortality selection might mean, at the same time that it has become far less specific about how those inequalities ought to behave. There is some demographic work (e.g., Manton et al. 1994, 1995; Woodbury and Manton 1986; Mohtashemi and Levins 2002) that has taken steps toward selection models that better capture our substantive understanding of cross-cutting, dynamic inequalities. This article hopes to inspire more.

REFERENCES

Dupre, M. E., Franzese, A. T. & Parrado, E. A. 2006. Religious attendance and mortality: Implications for the black-white mortality crossover. *Demography,* 43**,** 141-164.

Guillot, M. 2007. "Mortality in Kyrgyzstan since 1958: Real Patterns and Data Artifacts." Espace Populations Societies 2007(1):113-26.

Hill, M. E., Preston, S. H., and Rosenwaik, I. 2000. "Age Reporting Among White Americans Aged 85+: Results of A Record Linkage Study." Demography 37(2): 175-86.

Ho and Elo 2013. The Contribution of Smoking to Black-White Differences in U.S. Mortality. *Demography,* 50:2, 545-568.

Hoffman, F.L. 1896. The race traits and tendencies of the American Negro, *Publications of the American Economic Association* 11: 1–329.

Hoffmann, R. 2008. *Socioeconomic Differences in Old Age Mortality*. Springer: New York, NY.

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.

Huang, F., and Wu, C. 2010. "Old Age Mortality Crossover in China: Does the Public Health Insurance Matter?" *Population Research* 34(1):95-105 [in Chinese; abstract in English].

Kestenbaum, B. 1992. "A Description of the Extreme Aged Population Based on Improved Medicare Enrollment Data." Demography 29(4): 565-80.

Luo, Y., Hawkley, L. C., Waite, L. J., & Cacciopo, J. T. 2012. Loneliness, health, and mortality in old age: a national longitudinal study. *Social Science and Medicine*, 74(6): 907-14.

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.

Manton, K. G., Poss, S. S., & Wing. S. 1979. The Black/White Mortality Crossover: Investigation from the Components of the Perspective of Aging. *The Gerontologist*, 19(3): 291-300.

Manton, K. G. & Stallard, E. 1981. Methods for evaluating the heterogeneity of aging processes in human populations using vital statistics data: Explaining the black-white mortality crossover by a model of mortality selection. *Human Biology,* 53**,** 47-67.

Manton, K. G., Stallard, E., Woodbury, M. A. & Dowd, J. E. 1994. Time-varying covariates in models of human mortality and aging: Multidimensional generalizations of the Gompertz. *Journals of Gerontology,* 49**,** B169-B190.

Manton, K. G., Woodbury, M. A. & Stallard, E. 1995. Sex differences in human mortality and aging at late ages: The effect of mortality selection and state dynamics. *Gerontologist,* 35**,** 597-608.

Mohtashemi, M. & Levins, R. 2002. Qualitative analysis of the all-cause Black-White mortality crossover. *Bulletin of Mathematical Biology,* 64**,** 147-173.

Nam, C. B. 1995. "Another Look at Mortality Crossovers." *Social Biology* 42:133-43.

Olshansky, S.J. 1995. New Developments in Mortality. *The Gerontologist* 35(5):583-587.

Preston, S. 1976. Family sizes of children and family sizes of women. *Demography*, 13:1, 105-114.

Rogers, A. 1992. Heterogeneity and selection in multistate population analysis. *Demography,* 29**,** 31-38.

Sautter, J. M., Thomas, P. A., Dupre, M. E., & George, L. K. 2012. Socioeconomic status and the Black-White mortality crossover. *American Journal of Public Health*, 102(8): 1566-71.

Sibley, E. 1930. *Differential Mortality in Tennesse, 1917-1928*. Fisk University Press: Nashville, TN.

Steptoe, A., Aparna, S., Panayotes, D., & Wardle, J. 2013. Social isolation, loneliness, and all-cause mortality in older men and women. *Proceedings of the National Academy of Sciences*, 110:15, 5797–5801.

Thornton, R. 2004. The Navajo-U.S. Population Mortality Crossover Since the Mid-20th Century. *Population Research and Policy Review*, 23:3, 291-308.

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.

Woodbury, M. A. & Manton, K. G. 1983. A mathematical model of the physiological dynamics of aging and correlated mortality selection. 1: Theoretical development and critiques. *Journals of Gerontology,* 38**,** 398-405.

Yao, L. & Robert, S. A. 2011. Examining the Racial Crossover in Mortality between African American and White Older Adults: A Multilevel Survival Analysis of Race, Individual Socioeconomic Status, and Neighborhood Socioeconomic Context. *Journal of Aging Research*, 1-8.

Yi, Z., & Vaupel, J. W. Oldest-Old Mortality in China. *Demographic Research*, 8(7): 215-244.