



UNIVERSITY OF MINNESOTA

**Multidimensional Mortality Selection:
Why Individual Dimensions of Frailty Don't Act Like Frailty**

Elizabeth Wrigley-Field, PhD[†]
University of Minnesota

June 7, 2019

Working Paper No. 2019-02
DOI: <https://doi.org/10.18128/MPC2019-02>

[†]Address correspondence to Elizabeth Wrigley-Field, University of Minnesota, Institute for Social Research and Data Innovation, 50 Willey Hall, Minneapolis, MN 55417 or email ewf@umn.edu. Support for this work was provided by the Minnesota Population Center at the University of Minnesota (P2C HD041023).

Multidimensional Mortality Selection:

Why Individual Dimensions of Frailty Don't Act Like Frailty

Elizabeth Wrigley-Field
February 19, 2019 version

Abstract:

Theoretical models of mortality selection have great utility in explaining otherwise puzzling phenomena. The most famous example may be the black-white mortality crossover: at old ages, blacks outlive whites, presumably because few frail blacks survive to old ages while some frail whites do. Yet theoretical models of unidimensional heterogeneity, or frailty, do not speak to the most common empirical situation for mortality researchers—where some important population heterogeneity is observed while some is not. I show that, when one dimension of heterogeneity is observed and another is unobserved, neither the observed nor the unobserved dimension need behave like classic, unidimensional frailty models predict. For example, in a multidimensional model, mortality selection can increase the proportion of survivors that is disadvantaged, or “frail,” and can lead black survivors to be more frail than whites, along some dimensions of disadvantage. Transferring theoretical results about unidimensional heterogeneity to settings with both observed and unobserved heterogeneity produces misleading inferences about mortality disparities. The unusually flexible behavior of individual dimensions of multidimensional heterogeneity creates previously unrecognized challenges for empirically testing selection models of disparities, such as models of mortality crossovers.

The classical mortality selection model is a triumph of formal demography. It starts from the premise that people vary systematically in their risk of mortality and derives the conclusion that, as they age, cohorts are progressively reduced to a group of robust survivors. Models of mortality selection have been used to explain mortality phenomena such as mortality crossovers, reversals in the sign of a disparity (e.g., Berkman et al. 1989; Dupre et al. 2006; Eberstein et al. 2008; Fenelon 2013; Guillot 2007; Hoffman 2008; Huang and Wu 2010; Kestenbaum 1992; Lynch et al. 2003; Manton et al. 1979; Nam et al. 1978; Nam 1995; Pearl 1922; Rogers 2002; Thornton 2004; Thornton and Nam 1968; Zeng and Vaupel 2003); mortality deceleration, the slowing of mortality's rise with age (e.g., Beard 1959, 1971; Fukui et al. 1993; Horiuchi and Wilmoth 1997, 1998; Kannisto 1992; Lynch and Brown 2001; Lynch et al. 2003; Olshansky 1998; Thatcher et al. 1998; Vaupel et al. 1979; Vaupel and Yashin 1985); and mortality compression, the concentration of deaths into a small age range (e.g., Engelman et al. 2010; Kannisto 2000; Lynch and Brown 2001; Lynch et al. 2003).

But the classical mortality selection model does not speak to some of the most important questions in modern empirical mortality research, which concern the potential contribution of particular dimensions of heterogeneity in the context that other important dimensions are unobserved. The classical model is unidimensional: the heterogeneity that mortality selection acts on is captured by a single unobserved scalar fact about an individual, i.e., in disciplinary jargon, whether the individual is “frail” or “robust.”¹ This model, in both its binary and its continuous forms, developed historically in the context that old-age mortality data were limited, and creative theorizing made up for what could not yet be measured. Yet social science theories

¹ All standard models of mortality selection are unidimensional in this sense, whether frailty is modeled as binary (e.g., Lynch and Brown 2001, Lynch et al. 2003, Vaupel and Yashin 1985, Wrigley-Field 2014) or continuous (e.g., Gampe et al. 2010, Horiuchi and Wilmoth 1998, Missov and Finkelstein 2011, Vaupel et al. 1979).

of stratification are multidimensional and intersectional, and substantive knowledge of health stratification suggests that there are many overlapping, yet distinct, risk factors for mortality (Bowleg 2012). Increasingly, covariate-rich datasets allow some of these distinct dimensions of population heterogeneity to be measured and offer new opportunities to analyze the contributions to changing mortality disparities of particular heterogeneities, not just “frailty” as a black box. But since measured heterogeneity is always partial, work that foregrounds selection still needs to engage with unmeasured heterogeneity alongside measured covariates.

The need for a theory of multidimensional mortality selection is particularly apparent in recent empirical analyses of the black-white mortality crossover, which use unidimensional theory to ask multidimensional questions. The black-white mortality crossover is the phenomenon that black mortality exceeds white mortality at younger ages but falls below white mortality in old age, around age 85. The classical selection explanation for the crossover posits that blacks as a group are subject to greater selective pressure than whites, since their mortality is higher (e.g., Thornton and Nam 1968, Vaupel et al. 1979, Vaupel and Yashin 1985, Nam 1994, Lynch et al. 2003). Thus, old-age survivors include only the most robust members of the original black cohort, but a broader cross-section of the original white cohort, including frail members who would have been unlikely to survive had they been black. This longstanding theoretical explanation of the crossover has increasingly been engaged by empirical studies (Berkman et al. 1989, Dupre et al. 2006, Sautter et al. 2012, Yao and Robert 2011) that try to identify which particular, observed dimensions of heterogeneity might constitute pieces of this “frailty.” In current practice, research that aims to identify dimensions of heterogeneity that contribute to the crossover draws theoretically on mortality selection models designed to compare full populations (e.g., blacks vs. whites) in the presence of unobserved heterogeneity, but asks questions that rely

on complex, nested comparisons (e.g., blacks vs. whites with and without stratifying on a consequential health risk). This paper shows that the insights of the unidimensional selection model cannot be imported into the multidimensional setting and used to license the same kinds of predictions about the nested comparisons that apply to the full populations.

To address this gap between formal theory and empirical practice, I offer a model of the black-white crossover in the presence of multiple dimensions of heterogeneity and investigate its behavior. This work builds on extensive prior research into the behavior of covariates in survival models.² Yashin and Manton made several early advances in incorporating unobserved covariates into more empirically realistic survival analyses (such as those in which covariate measures are selectively missing for some observations [Yashin and Manton 1997]) and in estimating unobserved heterogeneity from survival models with an observed covariate and an assumed baseline distribution of the unobserved covariate (Yashin et al. 1985). More broadly, an early line of mortality selection research promised to meld the theoretical precision of mortality selection modeling with the empirical richness of new longitudinal data. This research tradition explored multidimensional models of mortality selection processes that focused on single-population phenomena, such as mortality deceleration (Manton et al. 1994, 1995; Manton and Woodbury 1983; Woodbury and Manton 1983), and more recently was picked up in theoretical work by Finkelstein and Esaulova (2008) and Finkelstein (2012). Finkelstein (2012) considers a two-dimensional frailty model in the context of mortality deceleration, and suggests an analytical approach of successively breaking populations into heterogeneous subpopulations defined by a

² The multidimensional mortality selection considered here differs from the multivariate mortality selection analyzed extensively elsewhere, as in “shared frailty models” (e.g., Henderson and Oman 1999, Guo and Rodriguez 1992, Vaupel 1988, Wienke 2010: 131-160). The former deals with multiple independent variables and the latter, multiple (correlated) survival-time outcomes.

single dimension of frailty, and then breaking those subpopulations into homogeneous groups, which is used in the analysis to follow. But mortality deceleration processes apply only to a single population analyzed as a whole, whereas analyzing mortality crossovers requires comparing the mortality selection processes unfolding in multiple populations (e.g., blacks and whites). Other analyses (Bretagnolle and Huber-Carol 1988, Henderson and Oman 1999; see discussion in Wienke 2010: 127-130) that, like the current paper, model multiple observed covariates in the presence of unobserved heterogeneity, focus on quantifying the bias in estimated covariate effects. Each of these strands of prior research forms the lineage for the current paper, which asks a different question: what happens to a mortality disparity—such as the black-white disparity—when we incorporate a new covariate that we hypothesize to be part of the mortality selection process? Does the disparity change in a predictable way? In particular, how does the presence, absence, or timing of a mortality crossover change when the “frailty” that produces the crossover is partially adjusted for?

In short, I analyze whether the insights developed from unidimensional mortality selection theory, in the context of the crossover, can be extended to incorporate covariates representing partial measures of population heterogeneity. I show that, in general, they cannot.

Multidimensional mortality selection and unidimensional mortality selection offer similar perspectives when all of the heterogeneity in a population is observed, or when none of it is observed. But unidimensional heterogeneity models offer no clear guidance about a multidimensional reality in which some dimensions of heterogeneity are observed, while others are unobserved. Yet this is the most common situation for social scientists studying mortality with datasets that include social and biological covariates representing some (but not all) of the heterogeneity within each population. The fact that individual dimensions of “frailty” need not

behave like frailty as a whole implies that, when selection is occurring along multiple dimensions simultaneously, one can't recover how it occurs along any one dimension (even qualitatively) without accounting for the other dimensions. This also implies that stratifying the crossover on observed heterogeneity offers quite limited information about the underlying selection processes. Nevertheless, I also show that it is possible to make some predictions about how the age at crossover responds to stratifying on key dimensions of heterogeneity, if certain key assumptions can be made. This provides a direction for developing more specific multidimensional selection theory in the context of mortality disparities.

I proceed by first presenting the core features of unidimensional mortality selection and then contrasting it with multidimensional mortality selection. In introducing the multidimensional model, I present two (alternative) predictions about how conditioning on an observed dimension of heterogeneity, in the presence of unobserved heterogeneity, should move the age at crossover. I show that key facts about unidimensional heterogeneity do not hold for partially-observed multidimensional heterogeneity, highlighting some previously unrecognized theoretical possibilities, such as *frailty increases* (mortality selection can lead populations to become more frail as they age) and *frailty reversals* (mortality selection can lead black survivors to be more frail than white survivors), that result from multidimensional models being intrinsically interactive. The distinctive behavior of individual dimensions of multidimensional heterogeneity has the consequence that neither prediction about the age at crossover is supported: conditioning on partial measures of “frailty” has essentially unpredictable consequences for the crossover without far more specific assumptions about latent parameter values.

Throughout, I adhere to the following terminological conventions. I consider two *populations*: blacks and whites. Populations may be stratified by one or two dimensions of

heterogeneity, which may be unobserved or observed. The unobserved dimension of heterogeneity is always called *frailty* (in the unidimensional model) or *residual frailty* (in the multidimensional model), while the observed dimension of heterogeneity is called *exposure*. I call populations stratified by one dimension of heterogeneity *subpopulations*, e.g. the subpopulation of robust blacks or the subpopulation of exposed whites. I call populations stratified by two dimensions of heterogeneity *groups*, e.g. the group of exposed robust blacks, or the group of unexposed frail whites. All populations, subpopulations, and groups are analyzed as closed cohorts.

Mortality Selection with Unidimensional Heterogeneity

The classic model of mortality selection with unidimensional heterogeneity will serve as a baseline for the distinctive dynamics of mortality selection with multidimensional heterogeneity.

Unidimensional Mortality Selection Model

The classical mortality selection model (e.g., Vaupel et al. 1979, Vaupel and Yashin 1985) divides the black and white populations along a single dimension of heterogeneity, called *frailty*. Frailty may be analyzed as a binary or a continuous variable; here I use binary frailty, resulting in four internally homogenous subpopulations defined by race $k = \{b, w\}$ and frailty $j = \{f, r\}$. Frailty is unobserved. The subpopulations have proportional Gompertz hazards,

$$\mu_{k,j}(a) = \alpha_{k,j} e^{\beta a} \quad (1)$$

with shared slope $\beta > 0$ over age $a \geq 0$ and intercepts $\alpha_{k,j}$. The subpopulation-specific intercepts are defined as

$$\begin{aligned}
 \alpha_{w,r} &= \alpha > 0 \\
 \alpha_{b,r} &= b\alpha \\
 \alpha_{w,f} &= f\alpha \\
 \alpha_{b,f} &= bf\alpha
 \end{aligned} \tag{2}$$

Thus, conditional on frailty, black subpopulations have higher mortality than white subpopulations in proportion $b > 1$ (the *black mortality multiplier*); and, conditional on race, frail subpopulations have higher mortality than robust subpopulations in proportion $f > 1$ (the *frail mortality multiplier*).³

Aggregate mortality for the black and white populations is a weighted average of the mortalities of the frail and robust subpopulations within each race,

$$\bar{\mu}_k(a) = \pi_k(a) \cdot \mu_{k,f}(a) + (1 - \pi_k(a)) \cdot \mu_{k,r}(a) \tag{3}$$

where $0 \leq \pi_k(a) \leq 1$ is the proportion of race k that is frail and $1 - \pi_k(a)$ is the proportion of race k that is robust at age a . The proportion frail, in turn, is given by

³ In the classical mortality selection literature in the context of crossovers, disadvantage has been operationalized in different ways, including: greater mortality for blacks than whites at all levels of frailty, with black and white frailty equal at baseline (Vaupel et al. 1979), greater mortality for blacks than whites specifically among the frail (Vaupel et al. 1985), and greater mortality for blacks than whites among the frail and a larger initial proportion of frailty among blacks at baseline (Lynch et al. 2003). This paper's model is consistent with the Vaupel et al. (1979) approach, which offers the neatest fit with the general empirical preference for proportional hazard models (by assuming black disadvantage for all cohort members, not just the frail) and allows the paper to highlight how multidimensional selection produces flexible crossover results even without one important source of model flexibility (namely, differences in black and white initial frailty distributions).

$$\begin{aligned}
 \pi_k(a) &= \frac{\pi(0) \cdot S_{k,f}(a)}{\pi(0) \cdot S_{k,f}(a) + (1 - \pi(0)) \cdot S_{k,r}(a)} \\
 &= \frac{\frac{\pi(0)}{1 - \pi(0)}}{\frac{\pi(0)}{1 - \pi(0)} + \frac{S_{k,r}(a)}{S_{k,f}(a)}}
 \end{aligned} \tag{4}$$

where $\frac{\pi(0)}{1 - \pi(0)}$, a constant, is the ratio of frail to robust members of the population at baseline

(assumed to be the same among blacks and whites⁴) and $\frac{S_{k,r}(a)}{S_{k,f}(a)}$ is the ratio of robust to frail

survivors within each race at age a , where $S_{k,j}(a) = \exp\left(-\int_0^a \mu_{k,j}(u) du\right)$. Since the frail die more quickly than the robust, the survivorship ratio increases, and the proportion frail, $\pi_k(a)$, decreases monotonically with age. Since blacks always have higher mortality in each subpopulation, but not necessarily in the aggregate, the crossover is an example of Simpson's paradox (e.g., Hernán et al. 2011, Hutchinson et al. 2000).

Four Facts about this Unidimensional Heterogeneity Model of Racial Disparities

The unidimensional frailty model of the black-white mortality crossover just presented makes two key assumptions, from which two results follow.

First, by assumption, conditional on frailty, black mortality exceeds white mortality at every age, $\mu_{b,j}(a) > \mu_{w,j}(a)$.

⁴ I assume the same baseline distributions of frailty in the black and white populations in order to focus the analysis of mortality crossovers cleanly on the dynamics of mortality selection, rather than other potential sources of racial difference in mortality. The main substantive points do not depend on this assumption. Nonetheless, if black and white frailty composition differed at birth, some aspects of the presentation of the results would differ, as I remark below in footnote 8. The particular parameter values at which the patterns illustrated below occur would also change.

Second, by assumption, conditional on race, the frail have higher mortality than the robust at every age, $\mu_{k,f}(a) > \mu_{k,r}(a)$.

From the second assumption, it follows that, within each race, the proportion of survivors that is frail declines monotonically over age, $\pi_k(a+1) < \pi_k(a)$ (Vaupel and Yashin 1985).

From the first and second assumptions together, it follows that, if blacks and whites have the same proportion frail at baseline, blacks have a smaller proportion of frail survivors than do whites at every subsequent age, $\pi_b(a) < \pi_w(a)$, $a > 0$ (Vaupel et al. 1979). Thus, the racial difference in the frailty of survivors results from the interaction of the between-race disadvantage of blacks and the within-race disadvantage of the frail.

These four generalizations will provide a crucial point of comparison for the multidimensional model to follow.⁵

Two roles of unidimensional frailty

A visual representation of the relationships given in Equations (1)-(4) will serve as a basis of comparison for multidimensional heterogeneity, which is introduced below. The interaction between the disadvantage of blacks and the disadvantage of the frail is depicted visually in Figure 1a, which illustrates the functional relationships in population mortality with unidimensional heterogeneity. Figure 1a shows that the black mortality multiplier affects the

⁵ These two derivations from unidimensional crossover models are well known (e.g., Vaupel et al. 1979) and follow from the widely known fact that, in a proportional hazards context with one unmeasured (e.g., frailty) and one measured (e.g., race) covariate, the unmeasured covariate leads to an underestimation of the effect of the measured covariate (see, e.g., Aalen 1988, Henderson and Oman 1999, Hougaard et al. 1994). The second assumption is the defining assumption of fixed-frailty models (Finkelstein 2012), which have wide application beyond the crossover, while the first assumption is particular to crossover models. The great achievement of selection models of the crossover is to make this first assumption compatible with the existence of a crossover (Vaupel et al. 1979, Vaupel and Yashin 1985).

mortality of robust and frail blacks, while the frail mortality multiplier affects the mortality of frail blacks and whites. The mortality of frail and of robust blacks each affect the proportion of black survivors that are frail, and each of those three terms, in turn, affects aggregate black mortality; and likewise for whites. (Figure 1a omits the parameters shared across all subpopulations: α , β , and $\pi(0)$.)

Figure 1b zooms in on the part of Figure 1a depicting the effects of the frailty mortality multiplier, f , on aggregate race-specific mortality, $\bar{\mu}_k(a)$, with +/- marks indicating the sign of each effect.⁶ It shows that f plays two competing roles in aggregate mortality. On one hand, increasing f raises aggregate mortality by raising the mortality of the frail, $f \xrightarrow{+} \mu_{k,f}(a) \xrightarrow{+} \bar{\mu}_k(a)$. On the other hand, because f raises the mortality of the frail at each age, it also lowers aggregate mortality by reducing the proportion of the frail that survive to old age, $f \xrightarrow{+} \mu_{k,f}(a) \xrightarrow{-} \pi_k(a) \xrightarrow{+} \bar{\mu}_k(a)$. The functional relationships shown qualitatively here are given quantitatively in Supplement 1.⁷ The interaction between these two roles of frailty create the potential for a mortality crossover.

⁶ The arrows in Figure 1 represent multiplicative effects; thus, the overall sign of a path is the product of the signs on each arrow.

⁷ Increasing the frailty multiplier can increase as well as decrease mortality in each racial population. The total effect of f on aggregate mortality in a population depends on which path dominates the other. Thus, there can be spans of ages at which population-level mortality would be lower with a larger frailty multiplier than with a smaller one, because a larger frailty multiplier means that fewer frail survivors remain.

A black-white crossover can occur regardless of the signs of the total effect of f on aggregate mortality in the black and white populations. When the total effect of frailty on mortality is less positive, or more negative, for the white population than for the black population at a given age, a crossover can occur. (Whether a crossover occurs additionally depends on whether the effect of frailty outweighs the black mortality disadvantage at the subpopulation level, as suggested by Figure 1a.)

The Black-White Mortality Crossover with Unidimensional Heterogeneity

A crossover occurs when aggregate white mortality exceeds aggregate black mortality,

$\bar{\mu}_b(a) - \bar{\mu}_w(a) < 0$. This crossover condition can be decomposed into three terms by rearranging

the expanded forms of black and white mortality as given in Equation 3:

$$\pi_b(a)(\mu_{b,f}(a) - \mu_{w,f}(a)) + (1 - \pi_b(a))(\mu_{b,r}(a) - \mu_{w,r}(a)) + (\pi_b(a) - \pi_w(a))(\mu_{w,f}(a) - \mu_{w,r}(a)) < 0 \quad (5)$$

The key point is that the four facts about mortality with unidimensional heterogeneity given above fully determine the sign of all three terms.

The first term is the black-white difference in the mortality of the frail, $\mu_{b,f}(a) - \mu_{w,f}(a)$, weighted by the proportion of the black population that is frail, $\pi_b(a)$. The second term is the black-white difference in the mortality of the robust, $\mu_{b,r}(a) - \mu_{w,r}(a)$, weighted by the proportion of the black population that is robust, $1 - \pi_b(a)$. These two terms are always positive because black mortality is always higher than white mortality, conditional on frailty.

The third term is the black-white difference in the proportion frail, $\pi_b(a) - \pi_w(a)$, weighted by the frail-robust difference in the mortality of whites, $\mu_{w,f}(a) - \mu_{w,r}(a)$. This term is always negative because its two factors have different signs: the black-white difference in the proportion frail is always negative, whereas the frail-robust mortality difference among whites is always positive.⁸ I call this third term the *frailty factor*. It represents the contribution of frailty-induced mortality selection to the racial difference in mortality. The frailty factor will illuminate the dynamics of the multidimensional heterogeneity model considered below.

⁸ That the black-white difference in the proportion frail is always negative depends on the assumption that blacks and whites are equally likely to be frail at birth. If blacks were more likely than whites to be frail at birth, then this term would become negative only if the greater selection against frailty among blacks over time outweighed the initial excess frailty among blacks.

Equation 5 highlights the tradeoff at the heart of the crossover in the unidimensional selection model: higher black disaggregated mortality, but lower frailty among black survivors. This makes crossover dynamics with unidimensional heterogeneity qualitatively simple: the only question is whether and when the white-black compositional difference will outweigh the black mortality disadvantage at the subpopulation level.

The first column of Table 1 summarizes the key features of mortality selection with unidimensional heterogeneity—the two roles of frailty, the four generalizations about unidimensional heterogeneity, and the sign-constrained decomposition of the crossover. These features will serve as a point of comparison with the multidimensional model that I introduce next.

Mortality Selection with Multidimensional Heterogeneity

The unidimensional mortality selection model is the central reference point for work on mortality crossovers. But it is the wrong reference point for recent empirical work on the black-white mortality crossover, which is fundamentally multidimensional. Recent studies on the crossover (Dupre et al. 2006, Sautter et al. 2012) ask what happens when a particular dimension of heterogeneity within black and white subpopulations is observed and other dimensions remain unobserved. These studies stratify on the observed dimension of heterogeneity and compare the ages at crossover between the resulting subpopulations and the age at crossover between the aggregate populations. To formalize the theory implicit in practice, I propose a model of mortality selection with partially-observed multidimensional heterogeneity, show that it behaves quite differently from mortality selection with unidimensional heterogeneity, and analyze the crossover age in the new model.

Multidimensional Mortality Selection Model

To demonstrate that multidimensional selection models with partially observed heterogeneity exhibit intrinsically different behaviors than the classical unidimensional selection model, I present a multidimensional model that differs from the classical unidimensional model in only one respect: each racial population is crosscut by not one, but two dimensions of fixed heterogeneity. The observed dimension of heterogeneity describes whether or not people suffered a deleterious *exposure* in utero (e.g. maternal smoking, which satisfies the model assumptions tolerably well: it raises mortality, is fixed at birth, and is relatively evenly distributed by race at birth, although white women do smoke more before and during pregnancy [Curtin and Mathews 2016]). The unobserved dimension of heterogeneity describes whether people are *residually frail* or residually robust.

The multidimensional model thus contains eight internally homogeneous groups defined by race $k = \{b, w\}$, observed exposure, $i = \{t, n\}$, and unobserved residual frailty, $j = \{f, r\}$. The groups have proportional Gompertz hazards,

$$\mu_{k,i,j}(a) = \alpha_{k,i,j} e^{\beta a} \quad (6)$$

with shared slope $\beta > 0$ over age $a \geq 0$ and group-specific intercepts $\alpha_{k,i,j}$. The intercepts are defined as

$$\begin{aligned}
 \alpha_{w,n,r} &= \alpha > 0 \\
 \alpha_{b,n,r} &= b\alpha \\
 \alpha_{w,n,f} &= f^*\alpha \\
 \alpha_{b,n,f} &= bf^*\alpha \\
 \alpha_{w,t,r} &= t\alpha \\
 \alpha_{b,t,r} &= bt\alpha \\
 \alpha_{w,t,f} &= tf^*\alpha \\
 \alpha_{b,t,f} &= btf^*\alpha
 \end{aligned} \tag{7}$$

where $b > 1$ is the black mortality multiplier, as before; $f^* > 1$ is the residual frailty mortality multiplier; and $t > 1$ is the exposure mortality multiplier. (The exposed groups are designated with t as in tobacco exposure, or treatment.) I assume that, at baseline, both unobserved residual frailty and observed exposure are distributed independently of race, though not necessarily of each other.

The group-specific mortalities with multidimensional heterogeneity are analogous to the subpopulation-specific mortalities with unidimensional heterogeneity. In the multidimensional model, each set of subpopulations defined by one dimension of heterogeneity, aggregating over the other dimension (e.g., tobacco-exposed whites, aggregated over residual frailty), is a separate instantiation of the unidimensional model.

If both dimensions of heterogeneity were observed, then the black and white populations could be analyzed straightforwardly in terms of their component groups. If neither dimension of heterogeneity were observed, then the black and white populations could be analyzed as having just one dimension of heterogeneity with four (rather than two) categories, i.e., as a version of the classical unidimensional heterogeneity model. The multidimensional selection model speaks to a third situation—which rests at the heart of recent empirical work on the crossover—where one dimension of heterogeneity is observed and the other is unobserved.

Mortality in the subpopulation defined by race k and observed exposure i , $\bar{\mu}_{k,i}(a)$, is the weighted average of the residually frail and residually robust groups in the subpopulation,

$$\bar{\mu}_{k,i}(a) = \pi_{k,i}(a) \cdot \mu_{k,i,f}(a) + (1 - \pi_{k,i}(a)) \cdot \mu_{k,i,r}(a) \quad (8)$$

where $\pi_{k,i}(a)$ is the proportion of frail members of the subpopulation with exposure i , and $1 - \pi_{k,i}(a)$ is the proportion robust.⁹ By assumption, $\bar{\mu}_{k,i}(a)$ is observed, but its component parts are not.

Aggregate mortality of race k , $\bar{\mu}_k(a)$, is a weighted average of the subpopulation-specific mortalities, $\bar{\mu}_{k,i}(a)$,

$$\bar{\mu}_k(a) = T_k(a) \cdot \bar{\mu}_{k,f}(a) + (1 - T_k(a)) \cdot \bar{\mu}_{k,r}(a) \quad (9)$$

All terms of Equation 9 are observed, and $T_k(a)$ is the proportion of each race that is exposed,

$$T_k(a) = \frac{T_k(0) \cdot \pi_{k,f}(0) \cdot S_{k,f,f}(a) + T_k(0) \cdot (1 - \pi_{k,f}(0)) \cdot S_{k,f,r}(a)}{\left[T_k(0) \cdot \pi_{k,f}(0) \cdot S_{k,f,f}(a) + T_k(0) \cdot (1 - \pi_{k,f}(0)) \cdot S_{k,f,r}(a) + (1 - T_k(0)) \cdot \pi_{k,r}(0) \cdot S_{k,r,f}(a) + (1 - T_k(0)) \cdot (1 - \pi_{k,r}(0)) \cdot S_{k,r,r}(a) \right]} \quad (10)$$

By assumption, $T_k(a)$ is observed, but its component parts are not. Note that $T_k(a)$ is defined at the population level, whereas $\pi_{k,i}(a)$ is defined at the subpopulation level.¹⁰ The interaction

⁹ The formula for $\pi_{k,i}(a)$ in the multidimensional model is analogous to the formula for $\pi_k(a)$ in the unidimensional model given in Equation 4, replacing the subpopulation-level survivorships $S_{k,i}(a)$ in Equation 4 with the corresponding group-level survivorships $S_{k,i,j}(a)$ for the i th (exposed or non-exposed) subpopulation.

¹⁰ In the multidimensional model, I use uppercase Greek letters for composition defined at the population level (the [observed] proportion of each racial population that is exposed, aggregated over residual frailty, $T_k(a)$, and the [unobserved] proportion of each racial population that is residually frail, $\Pi_k(a)$, aggregated over tobacco exposure) and lowercase Greek letters for composition defined at the subpopulation level (the [unobserved] proportion of each exposure

between these two dimensions of heterogeneity drives the distinctive behavior of heterogeneity at the aggregate population level (as in $T_k(a)$, $\Pi_k(a)$) compared to heterogeneity at the subpopulation level (as in $\pi_{k,i}(a)$, $\tau_{k,j}(a)$), which the following sections will elucidate.¹¹

How Stratifying on Partially-Observed Heterogeneity Might Change the Age at Crossover: Two Predictions

Identifying particular dimensions of heterogeneity that contribute to the aggregate black-white crossover requires some testable prediction about observable phenomena derived from a multidimensional selection model. A natural place to look for testable predictions is in the outcome that dominates research on mortality selection: the age at onset for some mortality selection artifact (e.g., Berkman et al. 1989; Dupre et al. 2006; Horiuchi and Wilmoth 1997, 1998; Lynch and Brown 2001; Lynch et al. 2003, Sautter et al. 2012, Yao and Robert 2011), in this case the crossover. Thus, to connect the multidimensional heterogeneity model to the questions asked in empirical research, consider the question: What happens to the extent of racial disparities in mortality—and what happens to the age at mortality crossover, if any—when black and white mortality are stratified on an observed dimension of heterogeneity (“uterine tobacco exposure”) while another dimension (“residual frailty”) goes unobserved?

subpopulation that is residually frail, $\pi_{k,i}(a)$, and the [unobserved] proportion of each residual frailty subpopulation that is tobacco-exposed, $\tau_{k,j}(a)$).

¹¹ One could instead decompose population-level mortality into the aggregate proportion of each racial population that is residually frail ($\Pi_k(a)$, unobserved) and the proportion of each of those subpopulations that is exposed ($\tau_{k,j}(a)$, unobserved). Regardless of how population-level mortality is decomposed, it reflects the distribution of each race along both dimensions of heterogeneity simultaneously.

Here I develop two alternative predictions about how the age at crossover responds to such stratification. One prediction is derived from empirical literature on the crossover and the other follows from generalizing the behavior of unidimensional frailty to individual dimensions of multidimensional heterogeneity. These predictions do not follow directly from the unidimensional model, which is silent on multidimensional applications. Rather, they represent alternative attempts to generalize that model's logic to the questions asked in empirical practice. In what follows, I will assess how these predictions fare in describing the behavior of the crossover under partial stratification.

Prediction 1. Recent empirical work on the black-white mortality crossover (Dupre et al. 2006, Sautter et al. 2012) proceeds in three steps. First, it presents a hypothesis that mortality selection within the black and white populations operates simultaneously on an observed dimension of heterogeneity (such as poverty, education, or religiosity) and on an unobserved dimension of heterogeneity, residual frailty.¹² Second, it offers predictions about the ages at crossover in the aggregate and in the subpopulations when the black and white populations are stratified by the observed dimension of heterogeneity. Third, it tests these predictions in empirical data, concluding that the observed dimension is (in the case of poverty and religiosity) or is not (in the case of low education) a dimension of the heterogeneity that produces the crossover in the aggregate.

The predictions offered in the literature have the same structure for each of the various observed dimensions of heterogeneity. Both Sautter et al. (2012) and Dupre et al. (2006) use the

¹² The dimensions of heterogeneity explored in the empirical literature are traits that—unlike “frailty”—are acquired and lost by individuals over time. This extension of the classic mortality selection models to time-varying dimensions of heterogeneity can introduce significant complications (see Manton et al. 1994, 1995; Rogers 1992; Woodbury and Manton 1983; Vaupel et al. 1988; and Wrigley-Field 2013) that are not considered either in those papers or in this one. Here, I focus solely on how fixed dimensions of heterogeneity interact in the selection process.

criterion that a trait is “[a source] of heterogeneity in individual frailty that contribute[s] to the Black-White mortality crossover” (Sautter et al. 2012:1566) if two regression coefficients on mortality are statistically significant: the trait interacted with age, and the trait interacted with race.¹³ They further seem to take this criterion as coextensive with the criterion that the observed trait is part of the “frailty” (i.e. multidimensional heterogeneity) if and only if conditioning on the trait changes the age at crossover (in some direction). This prediction is not derived from any formal model of multidimensional heterogeneity. The Dupre/Sautter criterion, then, is that one can test a model with partially observed, multidimensional heterogeneity by conditioning on the observed dimension and assessing whether the age at crossover changes.

Prediction 2. The formal models referenced in this empirical literature include only unidimensional heterogeneity. In translating that model into a multidimensional setting, one might also expect a more specific prediction to hold. If each dimension of multidimensional heterogeneity—such as uterine tobacco exposure and residual frailty, or low education and residual frailty—behaved like unidimensional frailty, then each dimension would have a predictable effect on black-white disparities at any age, and on the age at crossover. In the unidimensional heterogeneity model, as summarized above, the frail have higher mortality than the robust, and more surviving whites than blacks are frail. If the same facts extended to individual dimensions of multidimensional heterogeneity, then the tobacco-exposed would

¹³ This criterion is explicit in Dupre et al. (2006:146): “To investigate whether religious involvement operates as a source of heterogeneity, two conditions must be satisfied and are hypothesized separately. First, in accordance with prior research that shows that religious involvement is more protective for blacks, the following hypothesis must be true: the effect of religious involvement will have a greater impact among blacks on the risk of dying [...]. Thus, blacks who attend services weekly or more will have a larger reduction in mortality than whites. Second, to support the claim that religion contributes to why hazard rates invert, the effect of religion must vary with age.” In Sautter et al. (2012), this criterion is implicit, but undergirds the empirical analysis.

necessarily have higher mortality than the non-exposed, and more surviving whites would necessarily be exposed to tobacco than surviving blacks (since if they were black and exposed they wouldn't survive). In that case, tobacco exposure would necessarily raise aggregate white mortality relative to black. Stratifying on observed tobacco exposure would therefore raise black mortality relative to white, delaying the crossover to an older age. In short, if each dimension of heterogeneity behaved like unidimensional heterogeneity, the aggregate black and white populations would have to reach a crossover before the subpopulations did. This would constitute a testable prediction of the multidimensional heterogeneity model.

In what follows, I will show that neither the Dupre/Sautter prediction, nor this more specific prediction about crossover order, follows from the multidimensional heterogeneity model.

Unexpected Behaviors of Multidimensional Heterogeneity: The Four Key Facts About Unidimensional Heterogeneity Do Not Apply

In the model of mortality selection with unidimensional heterogeneity, I identified four key facts and a resulting decomposition for the black-white mortality crossover in which all terms had known sign. None of these generalizations extend to the individual dimensions of multidimensional heterogeneity. That is, in a multidimensional heterogeneity scenario where some, but not all, dimensions of heterogeneity are observed, neither the observed nor the unobserved dimensions necessarily behave like unidimensional frailty.

The distinctive behaviors of the multidimensional model include phenomena that I label subpopulation race crossovers, frailty crossovers, frailty increases, and frailty reversals. The first two possibilities are straightforward extensions of the unidimensional model to the

multidimensional context; the latter two are more surprising departures from unidimensional selection.

I. Subpopulation race crossovers—In the unidimensional model, conditional on frailty, j , black mortality is always higher than white mortality, $\mu_{b,j}(a) > \mu_{w,j}(a)$. By contrast, in the multidimensional model, the subpopulations can have their own race crossovers. Conditional on unobserved residual frailty, black mortality can be either higher or lower than white mortality, $\bar{\mu}_{b,j}(a) \geq \bar{\mu}_{w,j}(a)$, at any given age. Analogously, conditional on observed exposure, black mortality can be either higher or lower than white mortality, $\bar{\mu}_{b,i}(a) \geq \bar{\mu}_{w,i}(a)$. For example, Figure 2 illustrates a cohort in which, in the exposed subpopulation, black mortality is higher than white mortality before age 70 and after age 76, but lower than white mortality in between. Figure 2 and all following numerical illustrations come from a large universe of simulations described and analyzed in Supplement 3; the specific parameter values for all illustrative figures are given in Table S3.2.

These black-white *subpopulation crossovers* can occur because each subpopulation defined by stratifying on the observed dimension of heterogeneity instantiates the unidimensional heterogeneity model given in Equations 1-3.

II. Frailty crossovers—In the unidimensional model, within each race, frail mortality is always higher than robust mortality, $\mu_{k,f}(a) > \mu_{k,r}(a)$. In the multidimensional model, within each race, the residually frail subpopulation may have either higher or lower mortality than the residually robust subpopulation, $\mu_{k,f}(a) \geq \mu_{k,r}(a)$, and at any age. Similarly, the exposed subpopulation

may have either higher or lower mortality than the non-exposed subpopulation, $\mu_{k,i}(a) \geq \mu_{k,n}(a)$.

Frailty crossovers are exactly analogous to black-white subpopulation crossovers. Figure 3 shows a frailty crossover, for a cohort in which residually frail mortality falls below residually robust mortality for both blacks (ages 61-80) and whites (ages 77-96).

III and IV. Frailty increases and frailty reversals—In the unidimensional model, survivors are progressively less likely to be frail as the population ages, $\pi_k(a+1) < \pi_k(a)$ (the third fact about the unidimensional model). Furthermore, given equal baseline frailty across races, black survivors are always less likely than white survivors to be frail after baseline, $\pi_b(a) < \pi_w(a)$ (the fourth fact).

By contrast, in a multidimensional model, mortality selection can increase, as well as decrease, population-level residual frailty, $\Pi_k(a+1) \geq \Pi_k(a)$, or population-level exposure, $T_k(a+1) \geq T_k(a)$. I call this possibility a *frailty increase*. Furthermore, mortality selection can make black survivors more or less likely than white survivors to be residually frail, $\Pi_b(a) \geq \Pi_w(a)$, or more or less likely to be exposed, $T_b(a) \geq T_w(a)$. I call this possibility that black survivors become more disadvantaged than white survivors a *frailty reversal*. Frailty increases and frailty reversals violate the most important insights into mortality selection derived from the unidimensional model.¹⁴

The formal conditions for frailty reversals are given in Supplement 2, but the intuition is straightforward. Just as unidimensional mortality selection creates a negative association

¹⁴ It is well known that frailty can increase in populations in which individuals can newly acquire frailty during their lives (see Vaupel et al. [1988] for one systematic exploration of population dynamics that can result from such dynamic frailty). It is specifically in the context of frailty fixed in individuals that the frailty increases and frailty reversals illustrated here are deeply surprising.

between race and frailty among survivors, multidimensional mortality selection creates a negative association between tobacco exposure and residual frailty, within each race. This negative association can become so strong that selecting *against* one of those dimensions of heterogeneity becomes selecting *for* the other. This can lead the dimension being selected *for* to increase over age (a frailty increase), or—because this selection is stronger among blacks—to become more common among blacks than among whites (a frailty reversal). When this occurs, the dimension being selected for is always the one with a weaker effect on mortality, because the selection for that dimension is driven by the complex associations created by selection against the stronger dimension. Thus, blacks will always be more selected than whites along the stronger dimension of heterogeneity, but not necessarily along the other dimension.

To illustrate frailty increases and a frailty reversal, Figure 4 shows the proportions of black and white survivors that are residually frail in a simulated cohort. Frailty increases occur for blacks from ages 83-94, and for whites from ages 90-101. These frailty increases result from frailty crossovers such that, in the black and white populations at these respective ages, the residually frail have lower mortality than the residually robust. Mortality selection at these ages therefore makes each population more residually frail.

A frailty reversal occurs from ages 86 to 97, during which black survivors are more likely than white survivors to be residually frail, although they were less likely to be residually frail before. Frailty reversals result from the interaction between the two dimensions of heterogeneity. In this cohort, exposure raises mortality a great deal at the individual level, while residual frailty raises mortality much less, $t \gg f^*$. Consequently, both races, and especially blacks, are heavily selected against exposure. Furthermore, all subpopulations, and especially those who are exposed, are selected against residual frailty. But since comparatively fewer exposed blacks than

whites survive, selection against residual frailty occurs predominantly among whites. The interaction of the selection against exposure and selection against residual frailty results in blacks being less selected against residual frailty than whites for an 11-year span.

Frailty increases and frailty reversals underscore just how different the multidimensional selection model is from the unidimensional model. When there is only a single dimension of fixed heterogeneity that raises mortality, the two things we are certain of is that it declines monotonically over age and that, if blacks and whites start out with the same proportion frail at baseline, they end up with fewer frail at each subsequent age. Neither of these core generalizations necessarily extend to each fixed dimension of heterogeneity that raises mortality when there is more than one. In the next section, I show that the interaction between the two dimensions of heterogeneity that drives these distinctive possibilities stems from a distinctive third role of frailty that is unique to the multidimensional model.¹⁵

Three Roles of Frailty in Multidimensional Mortality Selection

¹⁵ Frailty reversals in particular dramatize that the various dimensions of heterogeneity involved in mortality selection interact flexibly when heterogeneity is multidimensional. In the unidimensional model, the fourth fact about unidimensional heterogeneity is a necessary consequence of the first two facts: since blacks have higher mortality than whites (conditional on frailty) and the frail have higher mortality than the robust (conditional on race), black survivors will always be more heavily selected against frailty than white survivors. But in the multidimensional model, the fourth fact can fail to hold even when the first two facts do hold; a frailty reversal can occur even in the absence of subpopulation crossovers and frailty crossovers. That is, mortality selection can produce a larger proportion of residually frail members in the black population than the white population, regardless of whether blacks in aggregate have higher or lower mortality than whites, and regardless of whether the residually frail in aggregate have higher or lower mortality than the residually robust. Frailty reversals can arise, instead, because the racial difference in population composition along one dimension of heterogeneity affects the racial difference in population composition along the other dimension (i.e., the size and sign of $\Pi_b(a) - \Pi_w(a)$ can alter the size—and even reverse the sign—of $T_b(a) - T_w(a)$, and vice-versa), which can give population composition a Simpson's paradox structure (as I show in Supplement 2). The new, third role of frailty discussed below not only forestalls the four key facts about unidimensional heterogeneity; it also breaks the dependencies between the facts.

The reason that the four key facts about unidimensional heterogeneity do not extend to each dimension of heterogeneity in the multidimensional model is that each dimension of heterogeneity in the multidimensional model plays three, rather than two, roles in determining population-level mortalities, $\bar{\mu}_k(a)$. Figure 5 visualizes, and Table 1 verbally summarizes, these three roles of heterogeneity, focusing on the three roles of unobserved residual frailty for convenience. (Analogous arguments apply to the three roles of observed exposure.) Figure 5 and Table 1 are based on the decomposition of population-level mortality given in Equation 9.¹⁶

In the unidimensional model, the mortality multiplier on being frail, f , played two roles in population-level mortality for each race: it simultaneously increased population-level mortality by increasing the mortality of the frail subpopulation, and it reduced population-level mortality by reducing the proportion of frail survivors. In the multidimensional model, unobserved residual frailty, j (and, by analogy, observed exposure, i), plays the same two roles. An increase in the residual frailty multiplier f^* increases population-level mortality by increasing the mortality of the residually frail within each subpopulation defined by levels of observed exposure,

$$f^* \xrightarrow{+} \mu_{k,i,f}(a) \xrightarrow{+} \bar{\mu}_{k,i}(a) \xrightarrow{+} \bar{\mu}_k(a);$$

and it decreases population-level mortality by decreasing the proportion frail within each subpopulation defined by observed exposure,

$$f^* \xrightarrow{-} \pi_{k,i}(a) \xrightarrow{+} \bar{\mu}_{k,i}(a) \xrightarrow{+} \bar{\mu}_k(a).$$

As in the unidimensional model, these two roles would suffice to produce a race crossover in aggregate mortalities.

¹⁶ Alternative decompositions of population-level mortality would yield a different set of “roles of frailty” without changing the main substantive point of this section, namely, that the two dimensions of heterogeneity interact in ways that can make population-level mortality dynamics unpredictable. The decomposition in Equation 9 was chosen to match empirical analyses because it represents mortality in terms of observable quantities. The proportion of survivors that is residually frail is represented at the subpopulation level, since it is observed only indirectly via its effect on the mortality of each subpopulation (Equation 8), which the proportion of survivors that is exposed is represented at the population level, where it is directly observable.

The third role of heterogeneity, by contrast, is new and considerably more complex: residual frailty in the multidimensional model affects population-level mortality by changing the proportion of the population that is exposed, $T_k(a)$. This means that the two dimensions of heterogeneity—unobserved residual frailty and observed exposure—interact. Even if the two dimensions of heterogeneity start out distributed independently of one another (at birth), they will become associated over age: survivors who are disadvantaged along one dimension are unlikely to also be disadvantaged along the other, since such multiply disadvantaged individuals are least likely to survive.¹⁷

The effect of residual frailty on population-level mortality via observed exposure composition is essentially unpredictable, for two reasons.

First, increasing the disadvantage associated with residual frailty, f^* , can either increase or decrease the proportion of survivors that are exposed, $T_k(a)$. Insofar as the disadvantage associated with residual frailty increases the mortality of the exposed subpopulation, $\bar{\mu}_{k,d}(a)$, it will *decrease* the proportion of survivors to subsequent ages that is exposed,

$f^* \xrightarrow{+/-} \bar{\mu}_{k,d}(a) \xrightarrow{-} T_k(a)$. Insofar as the disadvantage associated with residual frailty increases the mortality of the non-exposed subpopulation, $\bar{\mu}_{k,n}(a)$, it will *increase* the proportion of survivors

¹⁷ In the language of causal inference, the association occurs because mortality is a *collider* for its risk factors. Conditional on survival, those risk factors become associated. See Elwert and Winship (2015) for examples. The classical mortality selection model of the crossover can be expressed as a model in which mortality is a collider for race and frailty. In the multidimensional mortality selection model, mortality is a collider for race, observed exposure, and residual frailty, producing three-way associations between them over age, even if no association existed at baseline.

that is exposed, $f^* \xrightarrow{+/-} \bar{\mu}_{k,n}(a) \xrightarrow{+} T_k(a)$.¹⁸ When the total effect of the two paths from f^* into $T_k(a)$ is positive in one population at some age, the result can be a “frailty crossover” between the exposed and non-exposed and a “frailty” (i.e., observed exposure) increase in that population at that age. When the total effect of the two paths into $T_k(a)$ is larger among blacks than among whites for some span of ages, the result can be a “frailty” (i.e., observed exposure) reversal.¹⁹

Second, increasing the proportion of survivors that is exposed, $T_k(a)$, can either increase or decrease population-level mortality, $T_k(a) \xrightarrow{+/-} \bar{\mu}_k(a)$. Increasing $T_k(a)$ will increase population-level mortality when the exposed subpopulation has higher mortality than the non-exposed subpopulation. Increasing $T_k(a)$ will decrease population-level mortality when the exposed subpopulation has lower mortality than the non-exposed subpopulation, that is, when there has been a “frailty” (i.e., exposure) crossover. Thus, absent precise quantitative knowledge of the model parameters, the third role of residual frailty has an unpredictable effect on aggregate mortality, $f^* \xrightarrow{+/-} T_k(a) \xrightarrow{+/-} \bar{\mu}_k(a)$.²⁰

¹⁸ Whether the disadvantage associated with residual frailty, f^* , has a larger effect on the mortality of the non-exposed or the exposed (and whether these effects have the same sign) depends on whether the increased mortality of the residually frail groups outweighs the increased selection of the residually frail groups in each exposure subpopulation. Both effects are greater among the exposed, making the total effects of their competing signs unpredictable *a priori*.

¹⁹ A frailty reversal in observed exposure, in which black survivors are more likely than white survivors to be exposed for some span of ages, can occur when the total effect of the paths into T_k , *cumulative* over all prior ages, is larger for blacks than for whites.

²⁰ Whether the effects of covariates have *a priori* predictable or unpredictable sign is determined by the level of aggregation, not the dimension of heterogeneity. The mortality penalty associated with residual frailty, f^* , always has a negative effect on the proportion of survivors who are residually frail at the subpopulation level, $f^* \xrightarrow{-} \pi_{k,i}(a)$. But since the dimensions of heterogeneity interact at the population level, as illustrated in Figure 5, f^* can have either a negative or a positive effect on the proportion of survivors who are residually frail at the

In sum, in the multidimensional model, the various dimensions of heterogeneity within each population interact with each other, making it extremely difficult to relate any one observed dimension of heterogeneity to clean predictions about population-level mortality. If it is difficult to relate any given observed dimension of heterogeneity to aggregate mortality in any one population, then it is doubly difficult to relate an observed dimension of heterogeneity to mortality differentials between populations. Next I show that this has an important implication for empirical research: stratifying on an observed dimension of heterogeneity (while another dimension remains unobserved) can either increase or decrease the black-white disparity in mortality, and the resulting subpopulations can reach a crossover at either older or younger ages than the aggregate population.

Decomposition of the Aggregate Crossover with Multidimensional Heterogeneity: Conditioning on observed heterogeneity can move the age at crossover in either direction

Equation 11 decomposes the black-white crossover in aggregate mortality, $\bar{\mu}_b(a) - \bar{\mu}_w(a) < 0$, along the observed exposure dimension:

$$T_b(a)(\bar{\mu}_{b,I}(a) - \bar{\mu}_{w,I}(a)) + (1 - T_b(a))(\bar{\mu}_{b,N}(a) - \bar{\mu}_{w,N}(a)) + (T_b(a) - T_w(a))(\bar{\mu}_{w,I}(a) - \bar{\mu}_{w,N}(a)) < 0 \quad (11)$$

Equation 11 is exactly analogous to the decomposition of the black-white mortality disparity with unidimensional heterogeneity given in 5—except that, in the unidimensional case, the signs of all three terms were known *a priori* because they were determined by the key facts about unidimensional heterogeneity. By contrast, in Equation 11, those facts need not apply, and each of the three terms can be either positive or negative at any given age.

population level, $f^* \xrightarrow{+/-} \Pi_k(a)$. Hence, the effect of f^* on population mortality via its effect on the *population-level* proportion residually frail could be positive or negative, $f^* \xrightarrow{+/-} \Pi_k(a) \xrightarrow{+/-} \bar{\mu}_k(a)$.

The first two terms of Equation 11 are, respectively, the black-white difference in the mortality of the exposed, $\bar{\mu}_{b,e}(a) - \bar{\mu}_{w,e}(a)$, weighted by the proportion of the black population that is exposed, $T_b(a)$, and the black-white difference in the mortality of the non-exposed, $\bar{\mu}_{b,n}(a) - \bar{\mu}_{w,n}(a)$, weighted by the proportion of the black population that is non-exposed, $1 - T_b(a)$. These two terms can be either positive or negative because a black-white *subpopulation crossover* can either occur or fail to occur inside each subpopulation defined by observed exposure.

The third term of Equation 11 is the *frailty factor* representing the contribution of the racial compositional difference in observed exposure to the racial difference in aggregate mortality. It is the extent to which observed exposure is associated with higher mortality among whites, $\bar{\mu}_{w,e}(a) - \bar{\mu}_{w,n}(a)$, weighted by the black-white difference in observed exposure, $T_b(a) - T_w(a)$. This term can take either sign because each of its two factors can take either sign. The mortality difference will be positive, $\bar{\mu}_{w,e}(a) - \bar{\mu}_{w,n}(a) > 0$, as in the unidimensional case, as long as whites have not had a *frailty crossover* along the observed exposure dimension, and negative, $\bar{\mu}_{w,e}(a) - \bar{\mu}_{w,n}(a) < 0$, if they did have a frailty crossover. And the compositional difference will be negative, $T_b(a) - T_w(a) < 0$, as in the unidimensional case, as long as there has not been a *frailty reversal* along the observed exposure dimension, and positive, $T_b(a) - T_w(a) > 0$, if there has been a frailty reversal.

When the frailty factor is negative, the black-white mortality disparity is less positive, or more negative, in the aggregate than in the subpopulations. This can lead the aggregate to have a crossover when the subpopulations do not (as in the unidimensional model), or to have a more extreme crossover than the subpopulations do. Conversely, when the frailty factor is positive, the

black-white mortality disparity is more positive, or less negative, in the aggregate than in the subpopulations. This can lead to the absence of a crossover in aggregate mortality even when one of the subpopulations has a crossover.

Consequently, aggregate mortality can reach a crossover before both subpopulations, after both subpopulations, or in between the subpopulations. Stratifying black and white mortality on any single dimension of heterogeneity therefore moves the crossover in an essentially unpredictable direction. These results are summarized in the third panel of Table 1. Figure 6 shows illustrative examples, with solid vertical lines marking the onset of the aggregate crossover and dashed vertical lines marking the onset of the subpopulation crossovers.

Panel A of Figure 6 shows a cohort in which the aggregate populations reach a black-white mortality crossover before either of the subpopulations. Aggregate mortality starts to cross at age 80. The exposed subpopulation begins to cross at age 81, and the non-exposed subpopulation begins to cross at 89.

Panel B of Figure 6 shows a cohort in which the aggregate populations reach a black-white mortality crossover after the subpopulations. The exposed subpopulation reaches a crossover first, at age 56. The non-exposed subpopulation reaches a crossover next, at age 78, followed by the aggregate populations at age 90.

Aggregate mortality can cross in between the subpopulations as well. Panel C of Figure 6 shows a cohort in which the exposed subpopulations cross at age 78, the aggregate populations at age 80, and the non-exposed subpopulations at age 88.

Importantly, the crossover order—whether the aggregate populations reach a crossover at a younger or older age than the subpopulations—can change in response to very minor shifts in the model parameters. Cohorts that share most of their parameters can nevertheless vary in their

crossover order (shown in Supplement 3, which analyzes 1,635,687 simulated cohorts, and discussed below). Consequently, there is no obvious *a priori* prediction, absent strong assumptions about the latent model parameters, about whether stratifying on a single dimension of heterogeneity will increase or decrease the black-white mortality disparity at any given age, and whether it will increase or decrease the age at crossover in a cohort.

Implications for Empirical Research

These results cast doubt on the two potential tests of the multidimensional heterogeneity model based on stratifying black and white populations on an observed dimension of heterogeneity and comparing changes in the age at crossover to putative predictions based on the model. One potential test was based on the Dupre/Sautter prediction that stratifying on a single dimension of multidimensional heterogeneity should move the crossover in some (unspecified) direction. The results given above cast doubt on this criterion as either a necessary or a sufficient condition for identifying dimensions of heterogeneity that contribute to the aggregate crossover. First, the age at crossover will almost always shift in some direction when any trait is controlled for, as long as that trait is associated with both race and mortality. This is true regardless of whether that trait behaves like the *frailty* of a mortality selection model—that is, regardless of whether it meets (or tolerably approximates) the model assumptions of being fixed in individuals and raising mortality at all ages. Second, it is, however, possible for the crossover to occur at the same age in the aggregate population and in one of the subpopulations, even if the trait does constitute a dimension of frailty. Such a confluence of crossovers requires only—in the language of Equation 11—that the frailty factor has a very similar magnitude to the contribution of the other

subpopulation around the aggregate crossover age.²¹ Figure 7 shows an example of such a cohort. In this simulated cohort, the non-exposed subpopulation reaches a crossover at 18 days younger than the aggregate population—simultaneous ages from the perspective of any real study of old-age mortality.²²

The second potential test of the crossover model with an observed and an unobserved dimension of heterogeneity was based on the prediction that stratifying on the observed dimension might necessarily increase the black-white mortality disparity and delay the crossover. This would be true if individual dimensions of heterogeneity could be expected to behave like unidimensional heterogeneity. The results here cast doubt on this criterion as well. The preceding section shows that, in the multidimensional context, aggregate and subpopulation crossovers can in fact occur in any order. Thus, the strategy of empirically identifying particular dimensions of crossover-producing heterogeneity via such directional predictions similarly would not work.

The goal of identifying particular dimensions of heterogeneity that comprise a multidimensional analogue to “frailty” is an essential one for mortality research. But the results in this paper highlight the dangers of pursuing that goal without the benefit of an explicit model of multidimensional mortality selection. Moreover, they suggest that the goal may be surprisingly difficult to achieve. Demographic analyses of mortality commonly compare populations and subpopulations according to the age at which mortality selection artifacts

²¹ Simulations show that the aggregate can cross simultaneously with either the exposed or the non-exposed subpopulation. Simultaneous crossovers are defined as crossovers occurring at the same survivorship of the robust non-exposed whites (thus, the same age), to three decimal places. The simulation procedure is described in Supplement 3.

²² One might suspect that the aggregate crossover is nearly simultaneous with the non-exposed crossover because, by the time the aggregate crossover occurs, virtually all survivors are non-exposed. But this is not the case. At age 82, when the aggregate and non-exposed crossovers occur, 27% of black survivors and 24% of white survivors are exposed. (Thus, there has been a “frailty” reversal in observed exposure.)

begin—whether the crossover (Berkman et al. 1989, Dupre et al. 2006, Lynch et al. 2003, Sautter et al. 2012, Yao and Robert 2011) or mortality deceleration (e.g., Horiuchi and Wilmoth 1997, 1998; Lynch and Brown 2001; Lynch et al. 2003). But the flexibility of the multidimensional mortality selection model casts doubt on whether—absent much stronger assumptions about parameter values—any clear conclusions about the underlying heterogeneity can be drawn from the age at crossover in the aggregate population, compared to the subpopulations.

A question for further research is what other tests of the multidimensional mortality selection model might be possible. Supplement 3 shows that, while the crossover order varies with even small parameter changes over a large swath of simulated parameter space, some predictions nevertheless are possible, contingent on particular combinations of parameter values. The very presence of subpopulation crossovers implies that residual frailty is consequential and relatively common at baseline: whatever the measured exposure contributes to the aggregate crossover, the unmeasured heterogeneity is sufficient to generate a crossover.

In general, when the proportions of disadvantaged members (e.g., the residually frail and the exposed) are small at baseline, the crossover order is more constrained (largely because the aggregate dynamics will be dominated by the large group of more advantaged survivors); when the disadvantaged categories are larger at baseline, frequently, any order is possible even when the other parameters are fixed. Holding other parameters fixed, when baseline residual frailty is very high, it is relatively rare for the aggregate crossover to happen after both subpopulation crossovers when tobacco exposure is also very high at baseline, and it is relatively rare for the aggregate crossover to happen before both subpopulation crossovers when baseline tobacco exposure is low. (An aggregate crossover occurring between the two subpopulation crossovers is ubiquitous across the parameter space.)

The results here also imply that additional empirical tests of the multidimensional model may be possible in the special circumstance that the measured dimension of heterogeneity, t , can be assumed to represent a large portion of the total heterogeneity $f = t + f^*$. Such a scenario is presumably atypical in the case of covariates like religious participation (which likely account for only a relatively small part of the stable heterogeneity in mortality risk within racial populations), but might be reasonable in the case of a covariate like a Charlson Comorbidity Index (Charlson et al. 1994), which summarizes a variety of chronic medical conditions that collectively strongly predict mortality. These results suggest that a good strategy for empirical researchers might be to focus on covariates structured to capture much of the variation in mortality risk, such as by amalgamating many other covariates into a total measure of observed risk, rather than focusing on single covariates whose effects on mortality are not overwhelmingly large.²³ A covariate capturing much of the total heterogeneity licenses more predictions because it acts more like unidimensional heterogeneity. First, if $t > f^*$, then “frailty” reversals and frailty crossovers along the measured (t) dimension are impossible.²⁴ Measuring the proportion exposed over age in each race would therefore potentially allow this model to be falsified, given the assumption that $t > f^*$.²⁵ Unfortunately, given the more typical scenario that $t < f^*$ (i.e., unmeasured heterogeneity is more consequential for individual-level mortality than measured heterogeneity), the prediction that follows is about the unmeasured residual frailty dimension,

²³ Any such covariates would need to be studied in a setting in which, or operationalized such that, they are fixed in individuals. For example, chronic illnesses acquired by middle adult or early-elderly ages might be used as a strong predictor of mortality at older ages.

²⁴ Frailty reversals and increases might still occur along the residual frailty dimension, complicating the interpretation of the observed associations between tobacco exposure and mortality.

²⁵ Since frailty reversals and frailty crossovers will not always occur along the dimension of heterogeneity that less strongly increases mortality, only the presence, not the absence, of these phenomena constitute a test of this model.

and therefore not directly empirically testable. A second test may be possible even if $t < f^*$, if t is still “large.” Any crossover requires that the frailest white mortality exceed the most robust black mortality. If t and f^* are similar in magnitude, or more generally if t is large, then it is possible for $t + f^* = f > b$ while $f^* < b$ (even if $f^* > t$). In this situation, the observed subpopulations defined by exposure would never reach a crossover even though the aggregate population might—an empirically testable conclusion.²⁶

These empirical predictions—and the absence of similar predictions for measured covariates whose effect on mortality is small compared to the effect of the heterogeneity that remains unmeasured—suggest the value of explicit theorizing about mortality disparities in the presence of multiple dimensions of heterogeneity. They also suggest that, wherever possible, we attend to more localized parameter spaces, and that some of those spaces will be more revealing than others. In particular, to determine whether particular observed dimensions of heterogeneity contribute to a crossover through a multidimensional selection model, we should focus on dimensions that are highly consequential for mortality at the individual level. More generally, more specific predictions are possible as more assumptions are made to limit the simulation space to cohorts that better resemble U.S. cohorts (shown in Supplement 3). To the extent that this is a meaningful exercise in models that remain highly stylized, it suggests more room for developing fruitful predictions in the future.

Future work should explore whether other specifications of mortality disparities, such as alternative specifications of latent residual frailty (particularly gamma-distributed frailty [e.g., Gampe et al. 2010, Missov and Finkelstein 2011, Vaupel et al. 1979, Vaupel and Missov 2014]),

²⁶ Since subpopulation crossovers will not always occur even in subpopulations whose frailest whites have higher mortality than the most robust blacks, only the presence, not the absence, of subpopulation crossovers constitute a test of this model.

or its relationship to the hazard (Finkelstein and Esaulova 2006), and other ways of modeling underlying mortality inequalities besides proportional hazards (Steinsaltz and Wachter 2006), yield substantively similar results, and whether they might yield additional testable predictions. Another important avenue for future research is to attempt to quantify the bias that results in coefficient estimates for measured covariates when residual frailty is omitted from the model, extending work by Bretagnolle and Huber-Carol (1988) and Henderson and Oman (1999) to settings where the outcome is mortality disparities rather than the mortality or survivorship of single populations. Finally, the consequences of incorporating observed heterogeneity, in the presence of other, unobserved heterogeneity, should be explored alongside other avenues of making crossover modeling more substantively realistic. Among the possibilities are formally integrating baseline racial differences in exposure and residual frailty (Lynch et al. 2003) to the multidimensional selection process to look for further testable predictions given the range of parameter values²⁷; incorporating time-varying exposures that raise mortality risk (and hence selection) only in some spans of ages (e.g., Manton et al. 1994, Vaupel et al. 1988), which may open new ground for predictions based on comparing disparities in distinct age spans; and modeling the interactions between multiple observed traits in the presence of unobserved residual frailty. In particular, the evolving relationship between multiple observed dimensions of heterogeneity that are known to have large effects on individual-level mortality, in the presence of unobserved heterogeneity, is an important area in which to look for predictions that might let complex selection models be tested against empirical data.

²⁷ If blacks have higher baseline values of residual frailty or observed exposure, this will tend to delay a crossover in any units that aggregate over that respective dimension of heterogeneity, except in the presence of frailty crossovers.

Beyond mortality crossovers, these results have implications for any selection models of mortality disparities in the presence of partially-observed heterogeneity. Future research should also consider how these results about the associations created by selection between distinct dimensions of heterogeneity may affect analyses in other contexts, such as multi-decrement and multistate models. In a multi-decrement context, multidimensional heterogeneity opens opportunities to pursue models of multiple causes of death in which some heterogeneity raises the risk of every cause of death and some is specific to certain causes. Explicitly modeling multidimensional selection dynamics will be essential to interpreting such scenarios. Where such a structure of partially-shared risks can be assumed, looking separately at several causes of death may provide strategies to identify how observed and unobserved heterogeneity interact to produce specific outcomes.

In a multistate context, strategies for estimating unobserved heterogeneity may prove more challenging. Even with unidimensional heterogeneity, movement in and out of frailty can create a wide variety of selection patterns (e.g., Vaupel et al. 1988, Manton et al. 1994, Mohtashemi and Levins 2002, Wrigley-Field 2013). These patterns should also be explored in a multidimensional context. Since such dynamic models are less constrained than the model considered in this article, it seems likely that the aggregate patterns they can give rise to will be even more complex than those illustrated in the current analysis. If one (unobserved) dimension of heterogeneity is fixed and another (observed) dimension can be acquired during the life course, then, in order to interpret any mortality pattern, it becomes extremely important to model how the acquisition of the latter depends on the former.²⁸ Because the resulting dynamics are

²⁸ For example, imagine that the subpopulation with acquired “exposure” continually gains new members who, among the remaining survivors, are among the frailest along the fixed residual frailty dimension. The resulting associations between acquired exposure and fixed residual frailty are likely to be highly non-linear (as the subpopulation with acquired exposure is continually

complex, they may result in disparities (e.g., between blacks and whites conditional on acquired exposure) that grow and shrink over age. A fruitful question for future research is whether the observable size and hazards of the subpopulation with acquired exposure can be used to estimate the simultaneous processes of selection into exposure and into death. Empirical analyses that model mortality, and especially disparities in mortality, in the presence of acquired disadvantages alongside unobserved heterogeneity are likely to be very sensitive to model assumptions. Researchers exploring such analyses (e.g., Sautter et al. 2012) may want to carefully test the robustness of their results to alternative specifications.

Conclusion

In this paper, I analyzed the black-white mortality crossover in the presence of multiple dimensions of heterogeneity within each race. The crossover represents a concrete example through which to understand the frequent circumstance that one dimension of heterogeneity is observed and another is theorized to exist, but is unobserved. This situation is not captured by standard unidimensional heterogeneity models of mortality selection, but it is common in empirical research on the crossover (e.g., Berkman et al. 1979, Dupre et al. 2006, Sautter et al. 2012, Yao and Robert 2011), and in empirical studies of mortality broadly. This situation is also likely to become a larger part of mortality research as more datasets with rich covariates and sufficient coverage of old ages become available. I showed that the most basic facts about the unidimensional theoretical model do not necessarily extend to the situation in which some heterogeneity is observed and some is not. Neither the observed nor the unobserved dimensions

selected by mortality and simultaneously replenished from the frailest members of an increasingly robust, but shrinking in size, pool). These non-linearities will reflect complex interactions between the multidimensional mortality selection processes explored in this article and the dynamic frailty processes explored in earlier work using unidimensional heterogeneity models.

of heterogeneity necessarily behave as the classic, unidimensional models would predict: individual dimensions of frailty do not behave the same way as frailty in total.

The standard, unidimensional mortality selection model of the crossover is well summarized by four key facts: conditional on frailty, blacks have higher mortality than whites; conditional on race, the frail have higher mortality than the robust; within each race, frailty declines monotonically over age; and, given equal frailty at baseline age, black survivors to subsequent ages are less likely than white survivors to be frail. These facts about unidimensional heterogeneity together allow the crossover to occur and make its dynamics qualitatively simple.

The multidimensional model, analyzed with respect to a single observed dimension of heterogeneity, behaves very differently. None of the four key facts need hold: conditional on observed heterogeneity, blacks may have either higher or lower mortality than whites; conditional on race, the frail may have either higher or lower mortality than the robust; frailty can increase or decrease over age; and black survivors may be either more likely or less likely than white survivors to be frail. Generalizations that apply to heterogeneity as a whole—including the generalizations that form the foundation of a mortality selection account of the crossover—need not apply to each dimension of heterogeneity individually.

These possibilities arise because multidimensional mortality selection creates complex and essentially unpredictable—absent strong assumptions about parameter values—associations between the dimensions of heterogeneity. The four facts about unidimensional heterogeneity operate at the level of homogeneous subpopulations defined by race and frailty. But when heterogeneity is multidimensional, subpopulations defined by race and a single dimension of heterogeneity remain heterogeneous with respect to the other dimension of heterogeneity. These heterogeneous subpopulations change their composition over age, producing qualitatively

complex behavior at the population level. Multidimensional mortality selection is complex because in it, the basic dynamic of unidimensional mortality selection occurs fractally, at many interacting levels simultaneously. The crossover with unidimensional heterogeneity is a straightforward example of Simpson's paradox, occurring at a single level. But the crossover with multidimensional heterogeneity gives rise to Simpson's paradoxes at several levels, so that the phenomena occurring at the surface level—population-level mortality—become less intuitive. Multidimensional populations are mixtures of unidimensional subpopulations, and the mixture need not behave like its ingredients.

These results have theoretical and practical consequences. Theoretically, they suggest that the intuitions derived from the longstanding tradition of unidimensional mortality selection theory do not apply to multidimensional mortality selection. The intuitions developed from unidimensional mortality selection theory work well when the frailty can be thought of as a cohesive whole, an amalgam of entirely unobserved traits. But when we want to move from that perspective to one in which we identify some individual components of the heterogeneity that produces the crossover—when we want to get specific about “frailty”—those intuitions can become deeply misleading. Even if dimensions of heterogeneity are independently distributed at birth, they become associated—with each other and with race—as a cohort ages due to their joint contribution to mortality. Any dimension of heterogeneity therefore carries information about all of the others. Conditioning on any single dimension is not merely conditioning on a noisy measure of overall heterogeneity—it is conditioning *selectively* on whatever dimension was observed. The consequences of selective stratification on an observed dimension of heterogeneity for population and subpopulation mortality can only be described with a model

that explicitly incorporates the joint distribution of each dimension of heterogeneity as it changes over age—whether the context is a mortality crossover, or any other mortality trajectory.

Practically, the interaction between the compositional changes in the observed and in the unobserved dimension of heterogeneity produces unpredictable crossovers in the resulting subpopulations and aggregate populations. Stratifying on partial measures of heterogeneity can either increase or decrease the black-white mortality disparity, and it can move the black-white crossover in the subpopulations to either older or younger ages compared to the aggregate crossover. Since the general form of the heterogeneity model does not generally constrain the crossover order in the aggregate population compared to the subpopulations, the crossover order is not an empirical confirmation or refutation of the general form of this model. The seemingly most natural way to refute or confirm the multidimensional model—following a long tradition of using the age at onset of mortality selection artifacts to understand selection processes (e.g., Berkman et al. 1989; Dupre et al. 2006; Horiuchi and Wilmoth 1997, 1998; Lynch and Brown 2001; Lynch et al. 2003, Sautter et al. 2012, Yao and Robert 2011)—may not work.

This article’s new results about multidimensional heterogeneity were cast in terms of the black-white mortality crossover because it offers a specific empirical context. Historically, many theoretical results about unidimensional mortality selection were developed in dialogue with this empirical context, and recent empirical papers on the crossover clearly demonstrate the need for multidimensional selection theory. However, the implications of the results shown here extend well beyond that context. This paper joins several others in analyzing the consequences of selection along multiple dimensions simultaneously (Bretagnolle and Huber-Carol 1988, Manton et al. 1995, Henderson and Oman 1999, Finkelstein 2012). The results that are new in this paper specifically concern the behavior of mortality disparities in a common research situation: when

mortality is stratified by partial measures of heterogeneity. Collectively, this research shows the need to carefully consider which theoretical results about frailty as a whole do or do not extend to partially-observed heterogeneity.

The gap between the theoretical work on mortality selection and the recent empirical work on the crossover echoes a wider divergence between two traditions of demographic research in which the study of the crossover has traditionally been situated. Classical demography was an intellectually distinctive field that produced a series of models that excel at shifting perspectives between population aggregates and the individual-level status transitions that produce them. These models are able to reveal a great deal about population processes even in the absence of rich data; indeed, the striking creativity of formal demography in this era was presumably spurred by the need to wring as much information as possible from the limited data of the time. The classic mortality selection models, which interpret population-level mortality patterns as the consequence of theorized, unobserved subpopulations defined by frailty, are very much of this tradition.

In contrast, much recent empirical work in demography can be characterized as part of a broader tradition of population studies, drawing inspiration from much of the social sciences. One consequence of this disciplinary broadening is greater substantive engagement with processes of social stratification. Inequalities in lifespan that cross-cut race are surely multiple and intersecting, not unidimensional. And the advent of richer datasets allows some of these multiple heterogeneities to be measured. Recent work on the crossover has tried to break open the black box of “frailty” by asking how particular observed dimensions of heterogeneity might interact in a selection model with other dimensions of heterogeneity that remain unobserved. But

it has asked these questions without the benefit of any formal model of the multidimensional selection process.

Recent mortality crossover research sits uneasily between these traditions of formal and empirical demography. One motivation for an explicitly multidimensional model of mortality crossovers is to attempt to unite these two demographic traditions so that the substantively realistic and interesting questions of the recent empirical literature can be addressed with formal precision. Unidimensional frailty models are elegant and powerful tools for answering unidimensional questions, but multidimensional research questions need multidimensional theory. In particular, even as datasets grow ever richer, frailty remains an essential concept for mortality studies, as long as the heterogeneity that we do not measure remains as consequential as what we do. The essential insight of all selection models—that observed associations, taken at face value, can mislead us about issues as fundamental as whose elderly years are spent in greater disadvantage, blacks or whites—remains as powerful and necessary as ever.

But knowing that selection against *something* must wholly or partially account for the crossover is only somewhat satisfying. Ultimately, we want to build theories about what frailty consists of, and to test those theories. The results here are a step toward that goal, albeit a more halting step than we might wish, since they suggest that some plausible avenues of testing selection models are fraught with difficulty. Compared to unidimensional theory, multidimensional theory will be far more sensitive to population parameters, including unobservable parameters, in ways that make observable outcomes hard to predict. As models are made more substantively realistic by incorporating more than one dimension of heterogeneity in mortality risk, if the age at crossover is to remain a useful metric for testing a mortality selection model and gauging its properties, then those tests will need to be based in far more specific,

substantively grounded—and fallible—assumptions about the unmeasured inequalities inside populations.

REFERENCES

- Aalen, O. O. 1988. "Heterogeneity in survival analysis." *Statistics in Medicine* 7(11).
- 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.
- Berkman, L., Singer, B., & Manton, K. (1989). Black/white differences in health status and mortality among the elderly. *Demography*, 26(4): 661-678.
- Bretagnolle J & Huber-Carol C. (1988) Effects of omitting covariates in Cox's model for survival data. *Scandinavian Journal of Statistics*, 15:125 –138.
- Bowleg, L. (2012). The Problem With the Phrase Women and Minorities: Intersectionality— an Important Theoretical Framework for Public Health. *American Journal of Public Health*, 102(7): 1267-1273.
- Charlson, M., Szatrowski, T. P., Peterson, J., & Gold, J. 1994. Validation of a Combined Comorbidity Index. *Journal of Clinical Epidemiology*, 47(11):1245-51.
- Curtin, S. C. & Mathews, T. J. (2016) Smoking Prevalence and Cessation Before and During Pregnancy: Data From the Birth Certificate, 2014. *National Vital Statistics Reports*, 65(1).
- 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.
- Elwert, F. & Winship, C. 2015. Endogenous Selection Bias: The Problem of Conditioning on a Collider Variable. *Annual Review of Sociology*, 40: 31-53.
- 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.
- Fenelon, A. 2013. "An examination of black/white differences in the rate of age-related mortality increase." *Demographic Research*, 29(17): 441-472.
- Finkelstein, M. 2012. "On ordered subpopulations and population mortality at advanced ages." *Theoretical Population Biology* 81: 292-299.
- Finkelstein, M. & Esaulova, V. 2006. "Asymptotic behavior of a general class of mixture failure rates." *The Advances in Applied Probability* 38 (1), 244–262.
- Finkelstein, M. & Esaulova, V. 2008. "On asymptotic failure rates in bivariate frailty competing risks models." *Statistics and Probability Letters* 78: 1174-1180.
- 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.

- Guillot, M. 2007. "Mortality in Kyrgyzstan since 1958: Real Patterns and Data Artifacts." *Espace Populations Societies* 2007(1):113-26.
- Guo, G. & Rodriguez, G. 1992. Estimating a Multivariate Proportional Hazards Model for Clustered Data Using the EM Algorithm, with an Application to Child Survival in Guatemala. *Journal of the American Statistical Association*, 87(420): 969-976.
- Henderson R & Oman P. (1999). Effect of frailty on marginal regression estimates in survival analysis. *Journal of the Royal Statistical Society, Series B*, 61:367– 379.
- Hernán, M. A., Clayton, D., & Keiding, N. 2011. The Simpson's paradox unraveled. *International Journal of Epidemiology*, 40(3): 780-5.
- 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.
- Hougaard, P., Mygelaard, P., & Borch-Johnsen, K. 1994. "Heterogeneity Models of Disease Susceptibility, with Application to Diabetic Nephropathy." *Biometrics* 50(4): 1178-1188.
- Hutchinton, J. W., Kamakura, W. A., & Lynch, J. G. Jr. 2000. Unobserved Heterogeneity as an Alternative Explanation for "Reversal" Effects in Behavioral Research. *Journal of Consumer Research*, 27(3): 324-344.
- Kannisto, V. (1992). Frailty and survival. *Genus*, 47(3-4): 101-118.
- Kannisto, V. (2000). Measuring the compression of mortality. *Demographic Research*, 3(6).
- Kestenbaum, B. 1992. "A Description of the Extreme Aged Population Based on Improved Medicare Enrollment Data." *Demography* 29(4): 565-80.
- Luo, Y., Hawkey, L. C., Waite, L. J., & Caciopo, 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. (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., 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. 1983. A Mathematical Model of the Physiological Dynamics of Aging and Correlated Mortality Selection: II. Application to the Duke Longitudinal Study. *Journals of Gerontology*, 38(4): 406-413.
- 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.
- Olshansky, S. J. (1998). On the biodemography of aging: A review essay. *Population and Development Review*, 24, 381-393.
- 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.
- 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.
- Steinsaltz, D. R. & Wachter, K. W. 2006. "Understanding Mortality Rate Deceleration and Heterogeneity." *Mathematical Population Studies* 13(1): 19-37.
- Thatcher, A. R., Kannisto, V. & Vaupel, J. W. (1998). *The force of mortality at ages 80 to 120*. Odense, Denmark: Odense University Press.
- 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. 1988. Inherited frailty and longevity. *Demography* 25 (2), 277-287.
- 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., & Missov, T. 2014. Unobserved population heterogeneity: A review of formal relationships. *Demographic Research* 31: 659-686.
- 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., Yashin, A. I., & Manton, K., G. 1988. "Debilitation's Aftermath: Stochastic Process Models of Mortality." *Mathematical Population Studies* 1(1): 21-48.
- Wienke, A. 2010. *Frailty Models in Survival Analysis*. Chapman & Hall/CRC.
- Wrigley-Field, E. 2014. Mortality deceleration and mortality selection: Three unexpected implications of a simple model. *Demography*, 51: 51-71.
- Wrigley-Field, E. and Elwert, F. 2016. "Mortality Crossovers from Dynamic Subpopulation Reordering." In Schoen, R., ed., *Dynamic Demographic Analysis*, Springer: 177-199.

- 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(4): 398-405.
- Yashin, A. I. & Manton, K. G. 1997. Effects of Unobserved and Partially Observed Covariate Processes on System Failure: A Review of Models and Estimation Strategies. *Statistical Science*, 12(1): 20-34.
- Yashin, A. I., Manton, K. G., & Vaupel, J. W. 1985. Mortality and Aging in a Heterogeneous Population: A Stochastic Process Model with Observed and Unobserved Variables. *Theoretical Population Biology* 27: 154-175.
- 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.
- Zeng, Y., & Vaupel, J. W. 2003. Oldest-Old Mortality in China. *Demographic Research*, 8(7): 215-244.

TABLES

Unidimensional Heterogeneity	Multidimensional Heterogeneity ¹
<p><u>Two roles of frailty in population-level mortality</u></p> <p>I.</p> <ol style="list-style-type: none"> 1. Frailty increases subpopulation-level mortality (for the frail subpopulation). 2. Frailty decreases the proportion of survivors that are frail (in the population). 	<p><u>Three roles of residual frailty in population-level mortality²</u></p> <ol style="list-style-type: none"> 1. Residual frailty increases group-level mortality (for the two frail groups). 2. Residual frailty decreases the proportion of survivors that are residually frail (in the two exposure subpopulations). 3. Residual frailty can increase or decrease the proportion of survivors that are exposed (in the full population).
<p><u>Four Key Facts about the Unidimensional Heterogeneity Model of Racial Disparities</u></p> <p>II.</p> <ol style="list-style-type: none"> 1. Conditional on frailty, black mortality exceeds white mortality. 2. Conditional on race, frail mortality exceeds robust mortality. 3. The share of survivors who are frail decreases with age. 4. Black survivors are less likely than white survivors to be frail. 	<p><u>The Four Key Facts Need Not Apply</u></p> <ol style="list-style-type: none"> 1. Conditional on residual frailty, black mortality can be higher or lower than white mortality (<i>subpopulation crossovers</i> are possible). 2. Conditional on race, residually frail mortality can be higher or lower than residually robust mortality (<i>frailty crossovers</i> are possible). 3. The share of survivors who are residually frail may increase or decrease with age (<i>frailty increases</i> are possible). 4. Black survivors can be less likely or more likely than white survivors to be residually frail (<i>frailty reversals</i> are possible).
<p><u>Decomposition of population-level mortality crossover (Equation 5)</u></p> <p>III.</p> <p>All terms have known sign.</p> <p>Stratifying on frailty increases black-white mortality disparity.</p> <p>Stratifying on frailty removes the crossover.</p>	<p><u>Decomposition of population-level mortality crossover (Equation 11)</u></p> <p>All terms have unknown sign.</p> <p>Stratifying on residual frailty can increase or decrease black-white mortality disparity.</p> <p>Stratifying on residual frailty can make age at crossover older or younger.</p>

Table 1. Key comparisons between mortality selection with unidimensional and multidimensional heterogeneity

¹ Everything stated about unobserved residual frailty, with respect to the multidimensional model in the right column, also pertains to the observed exposure dimension of heterogeneity.

² The roles of residual frailty depend on how mortality is decomposed. This table uses the decomposition of mortality given in the text and in Figure 5, in which exposure composition is represented at the population level and residual frailty composition is represented at the subpopulation level, in order to match empirical situations.

FIGURES




Figure 1a. Functional relationships between the race and frailty multipliers and aggregate mortality in the unidimensional mortality selection model.




Figure 1b. Two roles of unidimensional frailty in aggregate mortality.

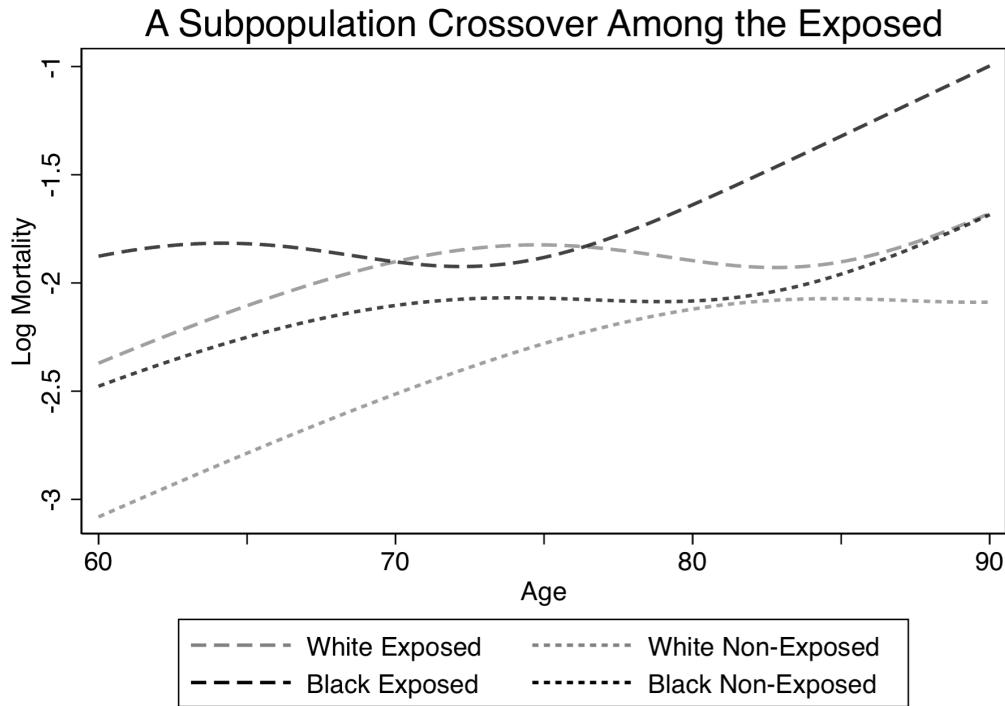


Figure 2. Subpopulation crossover: Black-white mortality crossovers can occur in the exposed subpopulation and in the non-exposed subpopulation.

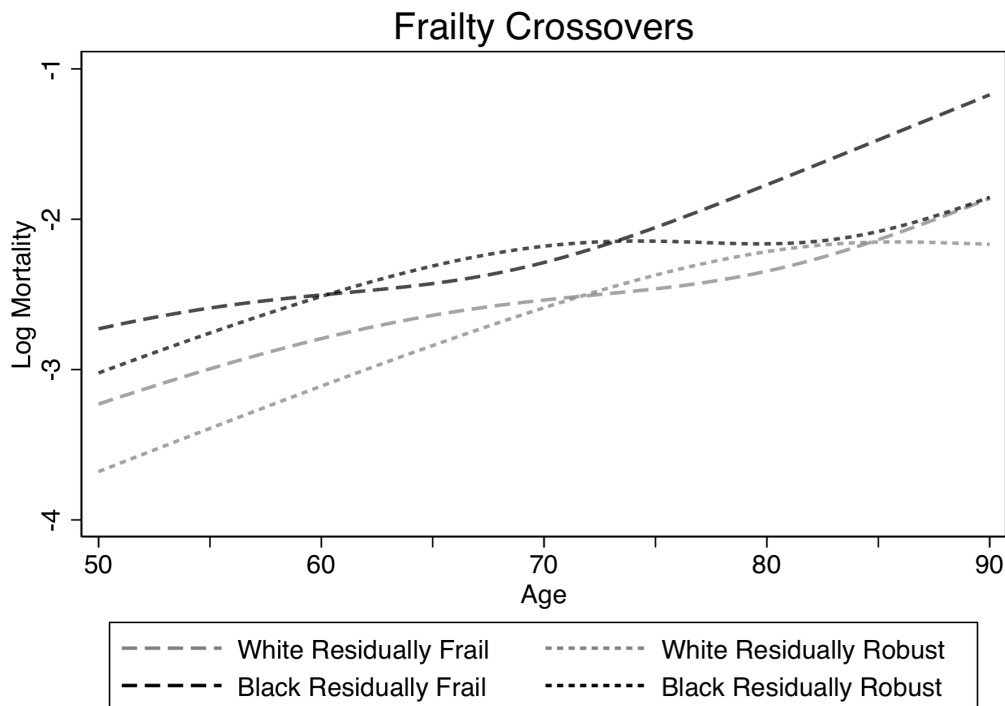


Figure 3. Frailty crossover: Conditional on race, residually frail members can have higher or lower mortality than residually robust members.

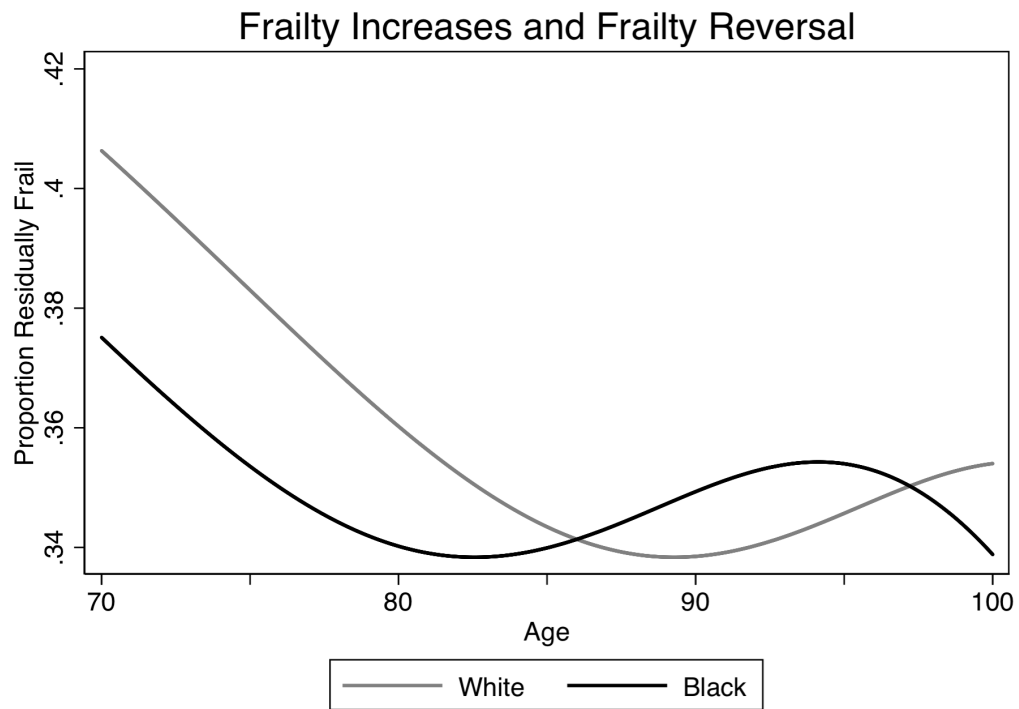


Figure 4. Frailty increases: The proportion of frail survivors increases among blacks and among whites. Frailty reversal: Whites can have a larger or smaller proportion of frail survivors over age.




Figure 5. Three roles of residual frailty on aggregate mortality in the two-dimensional mortality selection model.




Figure 6, Panel A. A simulated cohort in which black and white mortalities cross in the aggregate population before they cross in the exposed and non-exposed subpopulations.




Figure 6, Panel B. A simulated cohort in which black and white mortalities cross in the aggregate population after they cross in the exposed and non-exposed subpopulations.




Figure 6, Panel C. A simulated cohort in which black and white mortalities cross in the aggregate population after they cross in the exposed subpopulation and before they cross in the non-exposed subpopulation.

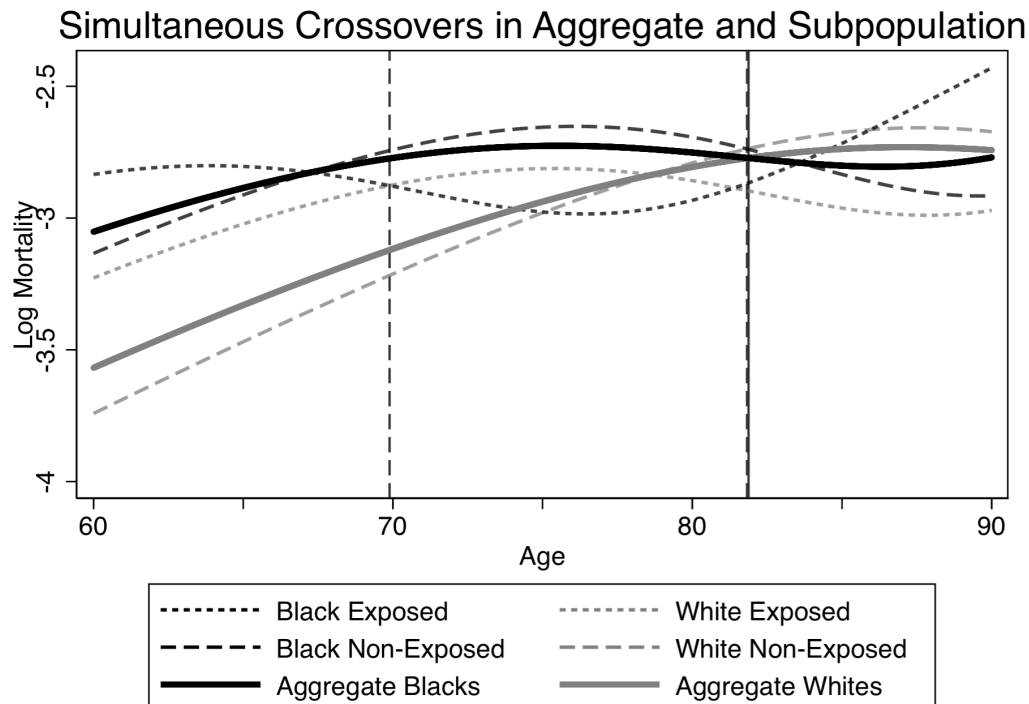


Figure 7. A simulated cohort in which aggregate mortality crosses essentially at the same time as the non-exposed (baseline) subpopulation. (These simultaneous crossovers begin just after the exposed subpopulation crossover ends.)

SUPPLEMENTARY APPENDIX 1:**Quantitative expressions for the roles of unidimensional and multidimensional heterogeneity in aggregate mortality**

Here I present partial derivatives quantifying the functional relationships between frailty/residual frailty and population mortality—the “roles of frailty”—illustrated qualitatively in Figure 1b and Figure 5 in the body of the paper.

Mortality selection with unidimensional heterogeneity

Equation S1.1 decomposes the contribution of the frailty mortality multiplier, f , to population-level mortality, $\bar{\mu}_k(a)$ for race k , into a contribution via subpopulation mortality, $\bar{\mu}_{k,j}(a)$, and a contribution via the proportion frail, $\pi_k(a)$.

$$\begin{aligned}
 \frac{\partial \bar{\mu}_k(a)}{\partial f} &= \frac{\partial}{\partial f} \left[\pi_k(a) \mu_{k,f}(a) + (1 - \pi_k(a)) \mu_{k,r}(a) \right] \\
 &= \frac{\partial}{\partial f} \left[\pi_k(a) (\mu_{k,f}(a) - \mu_{k,r}(a)) + \mu_{k,r}(a) \right] \\
 &= (\mu_{k,f}(a) - \mu_{k,r}(a)) \frac{\partial \pi_k(a)}{\partial f} + \pi_k(a) \frac{\partial \mu_{k,f}(a)}{\partial f}
 \end{aligned} \tag{S1.1}$$

Equation S1.1 says that a change in the frail mortality multiplier, f , can contribute to population-level mortality, $\bar{\mu}_k(a)$, via two pathways: the mortality of the frail, $\mu_{k,f}(a)$, and the proportion of the population that is frail, $\pi_k(a)$. (A third potential pathway, via the mortality of the robust, $\mu_{k,r}(a)$, is omitted from the third line of Equation S1.1 because the frailty multiplier is irrelevant

to the mortality of the robust, $\frac{\partial \mu_{k,r}}{\partial f} = 0$.) The remaining equations in this section quantify these two pathways, which represent the “two roles of frailty” in the unidimensional model.

In the first pathway, an increase in the frail mortality multiplier increases population-level mortality by increasing the mortality of the frail subpopulation, as shown in Equation S1.2:

$$\frac{\partial \mu_{k,f}(a)}{\partial f} = \mu_{k,r}(a) \quad (\text{S1.2})$$

A one-unit increase in the mortality multiplier on being frail, f , increases the mortality of the frail by the mortality of the robust, $\mu_{k,r}(a)$. And, as equation S1.1 showed, a one-unit increase in the mortality of the frail, $\mu_{k,f}(a)$, increases the mortality of the population by the proportion of the population that is frail, $\pi_k(a)$. The total contribution of the frail mortality multiplier to population-level mortality via its effect on the mortality of the frail is therefore the mortality of the robust weighted by the proportion of the population that is frail, $\pi_k(a) \cdot \mu_{k,r}(a)$. This contribution of frailty is always positive: the mortality multiplier f always increases population-level mortality with respect to its effect on the mortality of the frail.

In the second pathway, the frail mortality multiplier decreases population-level mortality by decreasing the proportion of the population that is frail, as shown in Equation S1.3:

$$\begin{aligned} \frac{\partial \pi_k(a)}{\partial f} &= \frac{\pi_k^2(a)(1 - \pi(0))S_{kr}(a)(\alpha_k - \mu_{kr}(a))}{\pi(0)S_{kf}(a)\beta} \\ &= \pi_k^2(a)\rho_k(a)\ln S_{kr}(a) \end{aligned} \quad (\text{S1.3})$$

A one-unit increase in the mortality multiplier on being frail, f , changes the proportion frail by

the amount $\pi_k^2(a)\rho_k(a)\ln S_{kr}(a)$, where $\rho_k(a) = \frac{(1 - \pi(0))S_{kr}(a)}{\pi(0)S_{kf}(a)}$, the ratio between, on the one

hand, the proportion robust at birth times the proportion of the original robust who survive to age

a , and, on the other hand, the proportion frail at birth times the proportion of the original frail who survive to age a . (There is no k subscript on the baseline proportions because the model assumes these do not vary by race.) And, as Equation S1.1 showed, a one-unit increase in the proportion frail, $\pi_k(a)$, increases the mortality of the population by the difference between frail and robust mortality, $\mu_{k,f}(a) - \mu_{k,r}(a)$.

The total contribution of the frail mortality multiplier to population-level mortality via its effect on the proportion of survivors that is frail is therefore the product of the difference between frail and robust mortality and the incremental change in the proportion frail, $(\mu_{k,f}(a) - \mu_{k,r}(a)) \cdot \pi_k^2(a) \rho_k(a) \ln S_{kr}(a)$. The crucial fact about this product is that it is never positive: all its terms are always positive except for $\ln S_{kr}(a) = \frac{\alpha_k - \mu_{kr}(a)}{\beta}$, which is always negative (or zero at age zero), since $\mu_{k,r}(a) > \alpha_k, a > 0$. Thus, the mortality multiplier f always decreases population-level mortality with respect to its effect on the proportion of the surviving population that is frail.

Equations S1.1-S1.3 are summarized in Equation S1.4, representing the total effect of frailty on the mortality of each population:

$$\begin{aligned} \frac{\partial \bar{\mu}_k(a)}{\partial f} &= \pi_k(a) \cdot \mu_{k,r}(a) + (\mu_{k,f}(a) - \mu_{k,r}(a)) \cdot \pi_k^2(a) \rho_k(a) \ln S_{kr}(a) \\ &= \pi_k(a) \mu_{k,r}(a) [1 + (f - 1) \pi_k(a) \rho_k(a) \ln S_{kr}(a)] \end{aligned} \quad (S1.4)$$

The sign of this expression can be positive or negative, and is determined by the sign of the bracketed term in the third line; its last term is always negative. The two roles of frailty expressed in Equations S1.2 and S1.3 play off against each other within each population, as summarized in Equation S1.4 and in Figure 1b.

Mortality selection with multidimensional heterogeneity

Equation S1.5 decomposes the contribution of the mortality multiplier on residual frailty, f^* , in the multidimensional heterogeneity model to population-level mortality, $\bar{\mu}_k(a)$ (for race k), into three pieces:

$$\begin{aligned}
 \frac{\partial \bar{\mu}_k(a)}{\partial f^*} &= \frac{\partial}{\partial f^*} \left[T_k(a) \pi_{k,f}(a) \mu_{k,f}(a) + T_k(a) (1 - \pi_{k,f}(a)) \mu_{k,f,r}(a) \right. \\
 &\quad \left. + (1 - T_k(a)) \pi_{k,n}(a) \mu_{k,n,f}(a) + (1 - T_k(a)) (1 - \pi_{k,n}(a)) \mu_{k,n,r}(a) \right] \\
 &= \frac{\partial \mu_{k,f}(a)}{\partial f^*} \cdot T_k(a) \pi_{k,f}(a) + \frac{\partial \mu_{k,n,f}(a)}{\partial f} \cdot (1 - T_k(a)) \pi_{k,n}(a) \\
 &\quad + \frac{\partial \pi_{k,f}(a)}{\partial f^*} \cdot T_k(a) (\mu_{k,f}(a) - \mu_{k,f,r}(a)) + \frac{\partial \pi_{k,n}(a)}{\partial f} \cdot (1 - T_k(a)) (\mu_{k,n,f}(a) - \mu_{k,n,r}(a)) \\
 &\quad + \frac{\partial T_k(a)}{\partial f^*} \cdot [\pi_{k,f}(a) (\mu_{k,f}(a) - \mu_{k,f,r}(a)) - \pi_{k,n}(a) (\mu_{k,n,f}(a) - \mu_{k,n,r}(a)) + \mu_{k,f,r}(a) - \mu_{k,n,r}(a)]
 \end{aligned} \tag{S1.5}$$

The residual frailty multiplier changes population-level mortality by changing the mortality of the residually frail (those who are and are not exposed, $\mu_{k,n,f}(a)$ and $\mu_{k,f}(a)$), the proportion of each population defined by exposure that is frail, $\pi_{k,n}(a)$ and $\pi_{k,f}(a)$, and the proportion of the population that is exposed, $T_k(a)$. (It does not change the mortality of the robust, $\mu_{k,n,r}(a)$ and $\mu_{k,f,r}(a)$, whether or not they are exposed; $\frac{\partial \mu_{k,n,r}}{\partial f} = 0$, $\frac{\partial \mu_{k,f,r}}{\partial f} = 0$, and those terms are omitted from Equation S1.5.)

The remainder of this appendix describes these three sets of pathways, which correspond to the “three roles of frailty” (more precisely, the three roles of each dimension of heterogeneity) in the multidimensional model. The key point is that the first two sets of pathways are analogous to the two pathways in the unidimensional heterogeneity model, but the third pathway, the effect of the residual frailty multiplier f^* on the proportion of the population that is exposed, $T_k(a)$, has no analogue in that model.

The first set of pathways represents the contribution of residual frailty, f^* to population-level mortality, $\bar{\mu}_k(a)$, via the direct effect of the mortality of the frail, $\mu_{k,n,f}(a)$ and $\mu_{k,d,f}(a)$.

This is given in Equation S1.6:

$$\frac{\partial \mu_{k,j,f}(a)}{\partial f} = \mu_{k,j,r}(a) \quad (\text{S1.6})$$

The effect of changing the residual frailty multiplier, f^* , on the mortality of each frail group, $\mu_{k,n,f}(a)$ and $\mu_{k,d,f}(a)$, is the mortality of the corresponding robust group, $\mu_{k,n,r}(a)$ and $\mu_{k,d,r}(a)$. As Equation S1.5 shows, these terms are weighted by how much of the total population is in each of these respective frail groups: they are multiplied by the respective proportion frail in each exposure subpopulation, $\pi_{k,n}(a)$ and $\pi_{k,d}(a)$, and the proportion of the total population that is in each exposure subpopulation, $1 - T_k(a)$ and $T_k(a)$. This is exactly analogous to Equation S1.2, corresponding to the unidimensional model, except for the additional weights representing the proportion exposed and non-exposed. As in Equation S1.2, the terms in Equation S1.6 are always positive: all else equal, the residual frailty mortality multiplier increases population mortality by increasing the mortality of the frail.

The second set of pathways represents the contribution of residual frailty, f^* , to population-level mortality, $\bar{\mu}_k(a)$, via the proportion residually frail in each subpopulation, $\pi_{k,n}(a)$ and $\pi_{k,d}(a)$. This is given in Equation S1.7:

$$\frac{\partial \pi_{k,j}(a)}{\partial f} = \pi_{k,j}^2(a) \rho_{k,j}(a) \ln S_{k,j,r}(a) \quad (\text{S1.7})$$

Since I have modeled the proportion residually frail as separate components of the residually frail among the non-exposed and among the exposed, $\pi_{k,n}(a)$ and $\pi_{k,d}(a)$, Equation S1.7 is exactly analogous to Equation S1.3, above, because each subpopulation in the multidimensional

model instantiates the unidimensional model. The effect of changing the residual frailty mortality multiplier, f^* , on the proportion of the subpopulation j that is frail, $\pi_{k,j}(a)$, is the same in the multidimensional model as in the unidimensional model. As shown in Equation S1.5, the total effect of changing the proportion frail in each subpopulation on population-level mortality,

$\bar{\mu}_k(a)$, comes from multiplying these $\frac{\partial \pi_{k,j}(a)}{\partial f}$ terms by the difference between frail and robust

mortality in the respective subpopulation, $\mu_{k,j,f}(a) - \mu_{k,j,r}(a)$, and the proportion of the population that is, respectively, non-exposed or exposed, $1 - T_k(a)$ and $T_k(a)$. Each of the resulting products is always negative: all else equal, the residually frail mortality multiplier reduces population mortality by reducing the proportion residually frail in each subpopulation.

The third pathway represents the contribution of residual frailty, f^* , to population-level mortality, $\bar{\mu}_k(a)$, via its effect on the proportion (aggregated over residual frailty) that is exposed, $T_k(a)$. This is given in Equation S1.8:

$$\frac{\partial T_k(a)}{\partial f^*} = \frac{\eta_{k,i,f}(a) \ln S_{k,i,r}(a) (\eta_{k,n,f}(a) + \eta_{k,n,r}(a)) - \eta_{k,n,f}(a) \ln S_{k,n,r}(a) (\eta_{k,i,f}(a) + \eta_{k,i,r}(a))}{\eta_{k,i,f}(a) + \eta_{k,i,r}(a) + \eta_{k,n,f}(a) + \eta_{k,n,r}(a)} \quad (S1.8)$$

Equation S1.8 shows that the mortality multiplier on residual frailty, f^* , alters the proportion of the total population that is exposed by altering the proportion of survivors in one of the exposed groups and one of the non-exposed groups (i.e., the respective groups that are residually frail), the effect of which on the population-level proportion exposed depends on the proportion of the population in each of the four groups within each race. The population proportions are given in terms of $\eta_{k,j,i}(a)$, the initial proportion in each group times its survivorship. For example,

$$\eta_{k,i,r}(a) = T(0)(1 - \pi_i(0))S_{k,i,r}(a).$$

The effect of the residual frailty mortality multiplier on the proportion exposed, $\frac{\partial T_k(a)}{\partial f^*}$, can be either positive or negative. Each term in the difference expressed in the numerator is negative (each is the product of a negative term and a positive term), while the denominator is positive. Thus, increasing the residual frailty mortality multiplier increases the proportion exposed when $\left| \eta_{k,f}(a) \ln S_{k,f,r}(a) (\eta_{k,n,f}(a) + \eta_{k,n,r}(a)) \right| < \left| \eta_{k,n,f}(a) \ln S_{k,n,r}(a) (\eta_{k,f}(a) + \eta_{k,f,r}(a)) \right|$, i.e., when the residual frailty mortality multiplier reduces the relative size of the exposed frail more than the relative size of the non-exposed frail, and decreases the proportion exposed when that inequality is reversed.

Finally, the effect of changing the proportion exposed was shown in Equation S1.5 to be $\pi_{k,f}(a)(\mu_{k,f}(a) - \mu_{k,f,r}(a)) - \pi_{k,n}(a)(\mu_{k,n,f}(a) - \mu_{k,n,r}(a)) + \mu_{k,f,r}(a) - \mu_{k,n,r}(a)$, which can itself take either sign. For example, increasing the proportion exposed can decrease population mortality when the proportion frail is much larger in the non-exposed subpopulation than the exposed subpopulation, $\pi_{k,n}(a) \gg \pi_{k,f}(a)$. (Changing the proportion exposed decreases mortality when exposure has had a *frailty crossover*, as defined in the text.) Thus, the residual frailty mortality multiplier can either increase or decrease the proportion exposed, and increasing the proportion exposed can either increase or decrease population mortality. The total effect of the residual frailty multiplier on population-level mortality via the proportion exposed,

$$\frac{\left[\pi_{k,f}(a)(\mu_{k,f}(a) - \mu_{k,f,r}(a)) - \pi_{k,n}(a)(\mu_{k,n,f}(a) - \mu_{k,n,r}(a)) + \mu_{k,f,r}(a) - \mu_{k,n,r}(a) \right] \cdot \left[\eta_{k,f}(a) \ln S_{k,f,r}(a) (\eta_{k,n,f}(a) + \eta_{k,n,r}(a)) - \eta_{k,n,f}(a) \ln S_{k,n,r}(a) (\eta_{k,f}(a) + \eta_{k,f,r}(a)) \right]}{\eta_{k,f}(a) + \eta_{k,f,r}(a) + \eta_{k,n,f}(a) + \eta_{k,n,r}(a)}, \text{ can take either}$$

sign.

The three roles of the mortality multiplier on residual frailty, f^* , on population mortality, $\bar{\mu}_k(a)$, under multidimensional heterogeneity are summarized in Equation S1.9:

$$\begin{aligned}
 \frac{\partial \bar{\mu}_k(a)}{\partial f} = & \mu_{k,f,r}(a) \cdot T_k(a) \pi_{k,f}(a) + \mu_{k,n,r}(a) \cdot (1 - T_k(a)) \pi_{k,n}(a) \\
 & + \pi_{k,f}^2(a) \rho_{k,f}(a) \ln S_{k,f,r}(a) \cdot T_k(a) (\mu_{k,f,f}(a) - \mu_{k,f,r}(a)) \\
 & + \pi_{k,n}^2(a) \rho_{k,n}(a) \ln S_{k,n,r}(a) \cdot (1 - T_k(a)) (\mu_{k,n,f}(a) - \mu_{k,n,r}(a)) \\
 & + \left[\pi_{k,f}(a) (\mu_{k,f,f}(a) - \mu_{k,f,r}(a)) - \pi_{k,n}(a) (\mu_{k,n,f}(a) - \mu_{k,n,r}(a)) + \mu_{k,f,r}(a) - \mu_{k,n,r}(a) \right] \\
 & \cdot \frac{\left[\eta_{k,f,f}(a) \ln S_{k,f,r}(a) (\eta_{k,n,f}(a) + \eta_{k,n,r}(a)) - \eta_{k,n,f}(a) \ln S_{k,n,r}(a) (\eta_{k,f,f}(a) + \eta_{k,f,r}(a)) \right]}{\eta_{k,f,f}(a) + \eta_{k,f,r}(a) + \eta_{k,n,f}(a) + \eta_{k,n,r}(a)}
 \end{aligned} \tag{S1.9}$$

This complex set of relationships, summarized visually in Figure 5, has an unknown total sign and undergirds the main results in this article.

SUPPLEMENTARY APPENDIX 2:

Frailty reversals

A *frailty reversal* along some dimension of heterogeneity (e.g., residual frailty) occurs in some span of ages when mortality selection causes the black population to have a larger proportion of disadvantaged survivors, along that dimension, than white survivors. Here I give the formal conditions for a frailty reversal to occur.

Equation S2.1 defines a frailty reversal in residual frailty at age a , such that the proportion of residually frail survivors is larger among blacks than among whites, $\Pi_b(a) > \Pi_w(a)$:

$$T_w(a) \cdot \pi_{w,f}(a) + (1 - T_w(a)) \cdot \pi_{w,n}(a) < T_b(a) \cdot \pi_{b,f}(a) + (1 - T_b(a)) \cdot \pi_{b,n}(a) \quad (\text{S2.1})$$

Aggregate residual frailty in each race, $\Pi_k(a)$, is a function of the proportion frail within each subpopulation defined by exposure, $\pi_{k,j}(a)$, weighted by the proportion of the race that is in each exposure subpopulation, $T_k(a)$, $1 - T_k(a)$.

Rearranging Equation S2.1 lets us decompose the conditions for a frailty reversal as the sum of three terms, as in Equation S2.2:

$$\begin{aligned} & T_b(a) \cdot (\pi_{w,f}(a) - \pi_{b,f}(a)) + (1 - T_b(a)) \cdot (\pi_{w,n}(a) - \pi_{b,n}(a)) \\ & + (T_b(a) - T_w(a)) \cdot (\pi_{w,n}(a) - \pi_{w,f}(a)) < 0 \end{aligned} \quad (\text{S2.2})$$

The left-hand side of Equation S2.2 has three terms: the white-black difference in the proportion frail among the exposed, $\pi_{w,f}(a) - \pi_{b,f}(a)$, weighted by the proportion of the black population that is exposed, $T_b(a)$; the white-black difference in the proportion frail among the non-exposed, $\pi_{w,n}(a) - \pi_{b,n}(a)$, weighted by the proportion of the black population that is non-exposed, $1 - T_b(a)$; and the black-white difference in exposure, $T_b(a) - T_w(a)$, weighted by the exposed - non-exposed difference in the proportion frail among whites, $\pi_{w,n}(a) - \pi_{w,f}(a)$.

A frailty reversal occurs when the sum of these three terms is negative. All terms in Equation S2.2 are necessarily non-negative except the black-white difference in aggregate exposure, $T_b(a) - T_w(a)$, which can take either sign but is positive as long as blacks and whites have *not* had a reversal in exposure. Thus, in a two-dimensional heterogeneity model, a reversal can occur only in a single dimension of heterogeneity (the one with the smaller mortality multiplier). The form of Equation S2.2 is analogous to the conditions for a black-white mortality crossover given in Equation 4 and Equation 10; a frailty reversal is, similarly, an example of Simpson's paradox. Moreover, the Simpson's paradox form implies that a frailty reversal can occur even in the absence of a frailty increase or a frailty crossover.

Substantively, Equation S2.2 states that a reversal in residual frailty requires: first, that whites have many exposed survivors compared to blacks; second, that few white exposed survivors are residually frailty compared to white non-exposed survivors; and third, that blacks and whites with the same exposure status have relatively similar proportions of residual frailty (especially in the tobacco subpopulation that blacks are most heavily clustered in). These conditions are sensible. In both the exposed and the non-exposed subpopulations, blacks have a smaller proportion of residually frail survivors than whites do. But in each race, the exposed also have a smaller proportion residually frail survivors than the non-exposed do. If blacks have a smaller proportion of exposed survivors than whites do, this can lead them to have more residually frail survivors: a frailty reversal.

In short, frailty reversals are driven by the three-way association between race, exposure, and residual frailty, and can occur only given a particular constellation of inequalities.

The conditions for a reversal in exposure, $T_b(a) > T_w(a)$, are exactly analogous and are given in Equation S2.3 for completeness:

$$\begin{aligned} & \Pi_b(a) \cdot (\tau_{w,f}(a) - \tau_{b,f}(a)) + (1 - \Pi_b(a)) \cdot (\tau_{w,r}(a) - \tau_{b,r}(a)) \\ & + (\Pi_b(a) - \Pi_w(a)) \cdot (\tau_{w,r}(a) - \tau_{w,f}(a)) < 0 \end{aligned} \tag{S2.3}$$

Such a black-white “frailty” reversal in exposure is the occurrence that can alter the sign of the third term in the main article’s Equation 10.

SUPPLEMENTARY APPENDIX 3:**Crossover Order in Simulated Cohorts with Multidimensional Heterogeneity**

Here I present the parameters for the examples shown in Figures 2-4 and 6-7 and discuss a wider set of simulation results. In particular, I show that the claim that the aggregate and the subpopulations can cross in any order is true not only over a large universe of simulation parameters, but also within many local areas of that simulation universe. In general, crossover order is unpredictable locally, not just globally; however, I highlight some notable exceptions.

The simulations add to the analysis in two ways. First, Equation 10 analyzed the crossover as a *state* and provided intuition for why population-level mortality can be crossed even when neither of the subpopulations are, or fail to be crossed even when both of the subpopulations are. But to fully speak to the age at crossover, one needs to view the crossover as an *event* that begins at a particular age. These perspectives can diverge if, for example, it were to turn out that the subpopulations can be crossed in the absence of the population-level crossover only when the population hazards have already crossed and uncrossed. The simulations verify that this is not the case: each of the crossover orders suggested in the body of the paper does actually occur.

Second, even if any crossover order is possible in some parameter space, it might turn out that crossover order is locally stable over a wide swath of parameter space. This would enable a multidimensional selection model, when supplemented by some relatively coarse assumptions about the parameter values, to generate a testable prediction about the order in which the population and the subpopulations reach a crossover. Simulating cohorts with a range of parameter values allows us to see whether such predictions are possible.

I choose a parameter space that is designed to maximize the number of crossovers while minimizing the assumptions about the sizes of the latent parameters. Accordingly, I use large values of the heterogeneity multipliers t and f and small values of the black mortality multiplier b to find crossovers, and relatively large values of the proportions in high-mortality groups at baseline so that those groups do not go extinct too quickly, but otherwise I simulate over a wide range of values. The parameters used in the unidimensional model and in the multidimensional model are summarized in Table S3.1, and the parameter values used in each example in the article are summarized in Table S3.2. The parameters for illustrative examples, unlike the systematic simulations, were chosen primarily with aesthetic criteria in mind (e.g., to make three distinct crossovers easily visible and distinguishable on each plot). The full set of parameter values explored is summarized in Table S3.3, and results in 3,422,250 simulated cohorts (i.e., every combination of these parameters) observed from birth to virtual extinction. Of these, 1,635,687 (48%) have at least one crossover, whether in the aggregate population or one of the subpopulations defined by tobacco exposure.

It is intrinsically difficult to calibrate parameters representing an amalgam of many different disadvantages, as “residual frailty” is imagined to be (and as “exposure” could be envisioned to be in the circumstance that many observed disadvantages are modeled together). The mortality multipliers on residual frailty and exposure used here represent levels of disadvantage found for demographic variables in extremely unequal conditions (e.g., sex differences in post-transition Russian mortality at old ages³¹), compound disadvantages (e.g.,

³¹ For Russian adults (aged 25-100) in cohorts born after 1950, the male/female ratio of age-specific mortality ranges from 2.8 to 4.8, averaging 3.7 (author’s calculations from Human Mortality Database data; accessed May 11, 2018 from <http://www.mortality.org/>).

Charlson Comorbidity Scores³²), or disadvantages that are relatively proximate to death³³, and are in line with values used in classic work on unidimensional mortality selection (such as Vaupel and Yashin 1985; see discussion in Wrigley-Field and Elwert 2016:196-198).

The results are fully generated by the mortality multipliers and the baseline frailty proportions; the baseline mortality intercept α and slope β , shared across groups, do not affect crossover order. This is because crossovers are fully defined when time is measured with reference to survivorship (e.g., with “age” measured as the proportion of the lowest-mortality group still alive, rather than as age in years), from which, in these proportional hazard models, α and β cancel out. Since age is a monotone (though non-linear) transformation of survivorship, the crossover order is identical in the age scale as in the survivorship scale. In order to limit the dimensionality of the simulations (for computational reasons and parsimony), I generate the simulations in survivorship scale, including all and only those parameters that affect crossover order, to create the full simulation universe.

I subsequently introduce α and β to create a more restricted, semi-realistic simulation universe, which I analyze alongside the full universe. To construct this universe, I employ a conservative measure of whether the cohorts generate aggregate Gompertz mortality parameters

³² In the original work validating this index, Charlson Comorbidity Scores amalgamate chronic disease burden and age, with each additional point reflecting either an additional comorbidity or a decade of age. A score of 4 (e.g., four comorbidities compared to zero, at the same age) is associated with a relative risk of mortality of 4.54 compared to a score of 0, a score of 5 carries a relative risk of 6.38, and a score of 6 carries a relative risk of 9.23 (Charlson et al. 1994:1249-50).

³³ For example, among patients with chronic obstructive pulmonary disease, the hazard ratio (adjusted for age and sex) associated with each standardized unit decrease in physical activity is 2.2; for a decrease in physical activity of two standard deviations, the hazard ratio is 4.7 (Waschki et al. 2011).

that resemble those of United States cohorts.³⁴ (The resulting cohorts are only “semi”-realistic in that the overall model remains highly stylized.) Specifically, I retain only those cohorts whose aggregate intercept is the same, to four digits, as those generated by some U.S. life table (1933-2015, calculated from Human Mortality Database data) and whose aggregate slope is within a range set by the slopes of the life tables in a five-year window surrounding that year.³⁵ For example, a cohort would be admitted if it has the same aggregate intercept as the U.S. life table in 1963, and an aggregate slope that is within the range of slopes in the 1961-1965 U.S. life tables.³⁶ This measure of realism is conservative in that the aggregate intercept and slope depend on the somewhat arbitrary additional assumption of an underlying α and β , and some cohorts (as defined by the mortality multipliers and baseline frailty proportions that fully determine the crossover order) might appear “realistic” with a different choice of α and β even though they do not appear realistic here. For each cohort, I use slopes of $\beta = .06$ and $\beta = .07$ (similar to the aggregate slope for recent cohorts³⁷) and three separate intercepts for each cohort that place the aggregate crossover at age 70, 80, or 90, respectively. A cohort is admitted to the restricted universe if it is “realistic” under any of the six combinations of these slope and intercept

³⁴ The Gompertz models are population-weighted, as if fit on individual-level mortality. To generate population-weighted model cohorts, I made blacks 17% of the population at baseline.

³⁵ I additionally admit cohorts whose intercept is either of two values representing “holes” in the U.S. life table intercept series, which declines monotonically over time: .0013 (the intercept is .0014 in 1936 and .0012 in 1937) and .0010 (the intercept is .0011 in 1938 and .0009 in 1939).

³⁶ By defining the slope range from adjacent years rather than the full series, this procedure implicitly accounts for the Strehler-Mildvan negative correlation between intercepts and slopes across populations (Strehler and Mildvan 1960). Over time, intercepts have fallen and slopes have risen.

³⁷ The Gompertz slope for the 2015 life table is .0725, up from .0465 in 1933. Aggregate slopes are smaller than underlying slopes, because of mortality selection.

parameters.³⁸ I further restrict this universe to the cohorts in which there is a crossover in the aggregate and in both subpopulations (which reduces the computational burden of estimating aggregate parameters for each cohort).

For all results below, I report results for the full universe of 1,635,587 simulated cohorts with at least one crossover, the universe limited to the 183,171 cohorts in which the aggregate and both subpopulations cross, and the universe further restricted to those 65,345 cohorts with realistic aggregate parameters.³⁹

Table S3.4 presents the crossover order in each simulation universe. In the full universe, the most common outcome is that the aggregate crosses before both subpopulations, as one might expect if generalizing from the unidimensional selection model.⁴⁰ However, when the universe is limited to those in which all three crossovers occur, the most common order is that the aggregate reaches a crossover after both subpopulations do. This discrepancy partly reflects that, in the full universe, 59% of the cohorts with a “first” crossover in the aggregate are actually

³⁸ On average, each cohort whose aggregate parameters are deemed ‘realistic’ is deemed realistic under four of the six combinations of intercept and slope. In all of the results reported here, I count each cohort once rather than weighting by the number of times each appears in the realistic universe.

³⁹ In the multidimensional model, aggregate black and white mortality can evince two distinct crossover intervals. I use the first crossover to define the crossover order.

⁴⁰ In the full universe of simulated cohorts, crossovers that do not occur are treated as occurring “later.” This follows the convention in empirical research, in which data limitations at very old ages leave it ambiguous as to whether or not crossovers that are not observed by the maximum ages with good data actually occur later. Thus, for example, in the full universe, I say that the aggregate crosses “first” if both subpopulation crossovers occur either at an older age or not at all.

Some of the 21% of cohorts with an indeterminate crossover order are cohorts in which the aggregate population crosses simultaneously with a subpopulation, but the vast majority are cohorts in which one subpopulation crosses while the other subpopulation and the aggregate do not. Thus, these indeterminate cohorts definitively are not aggregate-first cohorts, but are indeterminate between the aggregate crossing in the middle, last, or simultaneously with the later-crossing subpopulation.

crossovers in which only the aggregate populations ever reach a crossover at all. In the universe limited to cohorts that have all three crossovers and also generate realistic aggregate parameters, the most common outcome is a first crossover in the aggregate, but the outcomes are much more evenly split.

To analyze whether the crossover order is predictable within more localized parameter spaces, I divide the simulation parameters into two groups: mortality multipliers (e.g., excess mortality associated with being residually frail) and baseline distributions (e.g., proportion residually frail at baseline), and hold one set of parameters fixed while letting the others vary. This produces “cohort sets” that share all of their baseline mortality multipliers but vary in their baseline distributions, or the reverse. I analyze the crossover order within each cohort set.⁴¹ Locally valid predictions about crossover would require sets of cohorts in which all cohorts have the same crossover order.⁴²

I find that, for most values of baseline distributions and mortality multipliers, the resulting cohort sets contain multiple crossover orders. However, some local predictions are possible.

Figure S3.1 shows the crossover orders among sets of cohorts that share their mortality multipliers on being black, being exposed, and being residually frail, but vary in their baseline population composition (i.e., the baseline proportion exposed and proportion residually frail among the exposed and the non-exposed). The cohort sets are grouped by their mortality multipliers on being residually frail (y-axis) and being exposed (x-axis). In this and all following figures, the points are jittered to make more cohort sets visible. In order to collapse three

⁴¹ I ignore simultaneous crossovers.

⁴² For example: “When the mortality multipliers are set at x, y, z values, then the aggregate always crosses before the subpopulations, regardless of the baseline frailty distribution.” Or: “When the baseline distributions are set at x, y, z levels, then the aggregate always crosses in between the two subpopulations, regardless of the mortality multipliers.”

dimensions of variation into a two-dimensional graph, only cohort sets that have a black mortality multiplier equal to 2 are shown. The top panels show that, at low values of the mortality multiplier on residual frailty, the aggregate always crosses “first,” but that this largely reflects that the subpopulations are unlikely to cross at all. When the mortality multiplier for the exposed is very large and the mortality multiplier for residual frailty is only moderate, the aggregate crosses first or in the middle, but never last. But when both dimensions of frailty have a moderate or large effect on mortality, any crossover order is possible even among cohorts that share their mortality multipliers.

Figure S3.2 and Figure S3.3 show the crossover orders among sets of cohorts that share their baseline proportion residually frail among the exposed, baseline proportion residually frail among the non-exposed, and baseline proportion exposed. The cohort sets are grouped in Figure S3.2 by the baseline proportion exposed (x-axis, assumed to be observable) and the baseline proportion residually frail among the exposed (y-axis, assumed to be unobservable), and are shown for sets in which the baseline proportion residually frail among the non-exposed is .75. At low values of baseline frailty among the exposed, the aggregate never crosses last; this reflects that, in those cohorts, the exposed never reach a crossover. In cohorts in which all three crossovers occur, in general, when tobacco exposure is relatively rare at baseline, the aggregate crosses after or in between the subpopulations, and when tobacco exposure is relatively common at baseline, the aggregate crosses before or in between the subpopulations. This result makes sense: the more common tobacco exposure is, the more room there is for heightened mortality selection against tobacco exposure, among blacks compared to whites, to produce an aggregate crossover earlier than the crossovers driven solely by selection against residual frailty among each exposure subpopulation. At moderate levels of baseline residual frailty among the exposed

(among cohorts irrespective of their aggregate parameters) and at moderate levels of baseline tobacco exposure (among cohorts with realistic aggregate parameters), any crossover order is possible.

Figure S3.3 shows cohort sets grouped by their proportion exposed (x-axis, observable) and proportional residually frail among the non-exposed (y-axis, unobservable), among cohort sets for which baseline frailty among the exposed is .75. The results for the non-exposed are generally similar to those for the exposed, shown in the preceding figure⁴³, but seem to support more predictions: in the universe of cohorts that generate realistic aggregate parameters, no cohort set can take all three crossover orders. In these cohort sets, low baseline exposure implies that the aggregate will not cross first, and high baseline exposure implies that the aggregate will not cross last. However, this result turns on the procedure for choosing realistic cohorts being a good one. When the universe is not so restricted, then every part of the parameter space that has cohorts with all three crossovers has cohort sets that, despite sharing their baseline population composition, can generate every crossover order.

In general, the simulations simultaneously underscore the sensitivity of the crossover order to latent mortality parameters while also suggesting some avenues for identifying testable predictions about multidimensional selection models in the future. It is clear that neither the criterion used in earlier work by Dupre and Sautter, nor the criterion that seems to rise most naturally from unidimensional mortality selection, will work. But more localized predictions, including ones that attend to *whether* crossovers occur as well as when they do, may offer more promising avenues. In general, the more willing we are to limit our attention to only certain swaths of parameter space, whether because of the plausibility of the micro-level parameter

⁴³ Analogously to Figure S3.2, a major constraint is that, at low values of baseline residual frailty composition among the non-exposed, that subpopulation does not reach a crossover.

values entered into the model or the plausibility of the aggregate-level parameter values that result (or both), the more predictions are likely to be supported.

REFERENCES TO SUPPLEMENT 3

- Charlson, M., Szatrowski, T. P., Peterson, J., & Gold, J. 1994. Validation of a Combined Comorbidity Index. *Journal of Clinical Epidemiology*, 47(11):1245-51.
- Strehler, B.L., & Mildvan, A.S. 1960. General theory of mortality and aging. *Science* 132(3428): 14-21.
- Vaupel, J. W. & Yashin, A. I. 1985. Heterogeneity's ruses: Some surprising effects of selection on population dynamics. *American Statistician*, 39, 176-185.
- Waschki, B., Kirsten, A., Holz, O., Müller, K. C., Meyer, T., Watz, H., & Magnussen, H. 2011. Physical Activity Is the Strongest Predictor of All-Cause Mortality in Patients With COPD: A Prospective Cohort Study. *Chest*, 140(2): 331-342.
- Wrigley-Field, E. and Elwert, F. 2016. "Mortality Crossovers from Dynamic Subpopulation Reordering." In Schoen, R., ed., *Dynamic Demographic Analysis*, Springer: 177-199.

SUPPLEMENTARY APPENDIX 3 TABLES

Parameter	Model	Role
<i>Individual-level mortality parameters</i>		
β	UD,MD	Gompertz slope of mortality over age
α	UD,MD	Mortality intercept of lowest-mortality group
b	UD,MD	Mortality multiplier for blacks
f	UD	Mortality multiplier for frail
f^*	MD	Mortality multiplier for residually frail
t	MD	Mortality multiplier for exposed
<i>Compositional parameters</i>		
$\pi_k(a)$	UD	Proportion of racial subpopulation k that is frail
$\pi_{k,i}(a)$	MD	Proportion of subpopulation in race k and exposure i that is residually frail
$\Pi_k(a)$	MD	Proportion of race k that is residually frail, aggregated over exposure
$\tau_{k,j}(a)$	MD	Proportion of subpopulation in race k and residual-frailty j that is exposed
$T_k(a)$	MD	Proportion of race k that is exposed, aggregated over residual frailty
<i>Mortality</i>		
$\mu_{k,j}(a)$	UD	Mortality of subpopulation defined by race k and frailty j
$\bar{\mu}_k(a)$	UD	Mortality of race k , aggregated over frailty
$\mu_{k,j,i}(a)$	MD	Mortality of group defined by race k , residual frailty j , and exposure i
$\bar{\mu}_{k,i}(a)$	MD	Mortality of subpopulation defined by race k and exposure i , aggregated over residual frailty
$\bar{\mu}_k(a)$	MD	Mortality of race k , aggregated over residual frailty and exposure

Table S3.1. Parameters in the unidimensional and multidimensional mortality selection models.

Illustrated Phenomenon	Fig.	Parameters
Subpopulation Crossover	2	$\alpha = .00026569$, $\beta = .065$, $b = 2$, $f^* = 4$, $t = 2$, $\pi_{k,n}(0) = 0.9$, $\pi_{k,t}(0) = 0.95$, $T_k(0) = 0.5$
Frailty Crossover	3	$\alpha = .00034959$, $\beta = .06$, $b = 2$, $f^* = 2$, $t = 4$, $\tau_{k,r}(0) = 0.9$, $\tau_{k,f}(0) = 0.75$, $\Pi_k(0) = 0.95$
Frailty Increase and Frailty Reversal	4	$\alpha = .00006139$, $\beta = .06$, $b = 1.5$, $f^* = 2$, $t = 8$, $\tau_{k,r}(0) = 0.8$, $\tau_{k,f}(0) = 0.8$, $\Pi_k(0) = 0.5$
Aggregate mortality crosses before subpopulations	6A	$\alpha = .00006421$, $\beta = .06$, $b = 2$, $f^* = 8$, $t = 2$, $\pi_{k,n}(0) = 0.75$, $\pi_{k,t}(0) = 0.95$, $T_k(0) = 0.85$
Aggregate mortality crosses after subpopulations	6B	$\alpha = .00006007$, $\beta = .075$, $b = 2$, $f^* = 8$, $t = 4$, $\pi_{k,n}(0) = 0.95$, $\pi_{k,t}(0) = 0.75$, $T_k(0) = 0.75$
Aggregate mortality crosses between subpopulations	6C	$\alpha = .00014788$, $\beta = .06$, $b = 2$, $f^* = 6$, $t = 2$, $\pi_{k,n}(0) = 0.8$, $\pi_{k,t}(0) = 0.95$, $T_k(0) = 0.7$
Aggregate mortality crosses simultaneous with non-exposed subpopulation	7	$\alpha = .00010924$, $\beta = .06$, $b = 2$, $f^* = 8$, $t = 2$, $\pi_{k,n}(0) = 0.85$, $\pi_{k,t}(0) = 0.8$, $T_k(0) = 0.35$

Table S3.2. Parameter values for simulated cohorts in the text

Parameter	Range
b	1.05 and [1.25, 2.5] in units of .25
f^*	[2, 8] in units of .5
t	[2, 8] in units of .5
$\pi_n(0)$	[.25, .95] in units of .05
$\pi_t(0)$	[.25, .95] in units of .05
$T_0(a)$	[.25, .95] in units of .05

Table S3.3. Parameter values for simulation universe. Values of α and β are omitted in the full simulation universe because they do not affect crossover order. In the “realistic” universe, they are simulated at $\beta = .06$ and $\beta = .07$, and α values that place the aggregate crossover at ages 70, 80, and 90, respectively; this universe is then limited to cohorts whose resulting aggregate parameters resemble those of United States historical cohorts.

	Full universe (at least one crossover)	Universe with all three crossovers occurring	Universe with all three crossovers and realistic aggregate parameters
Aggregate first	52%	17%	42%
Aggregate middle	19%	25%	30%
Aggregate last	8%	55%	28%
Indeterminate order	21%	3%	0%
N	1,635,687	183,171	65,635

Table S3.4. Crossover order in simulated cohorts.

APPENDIX 3 FIGURES

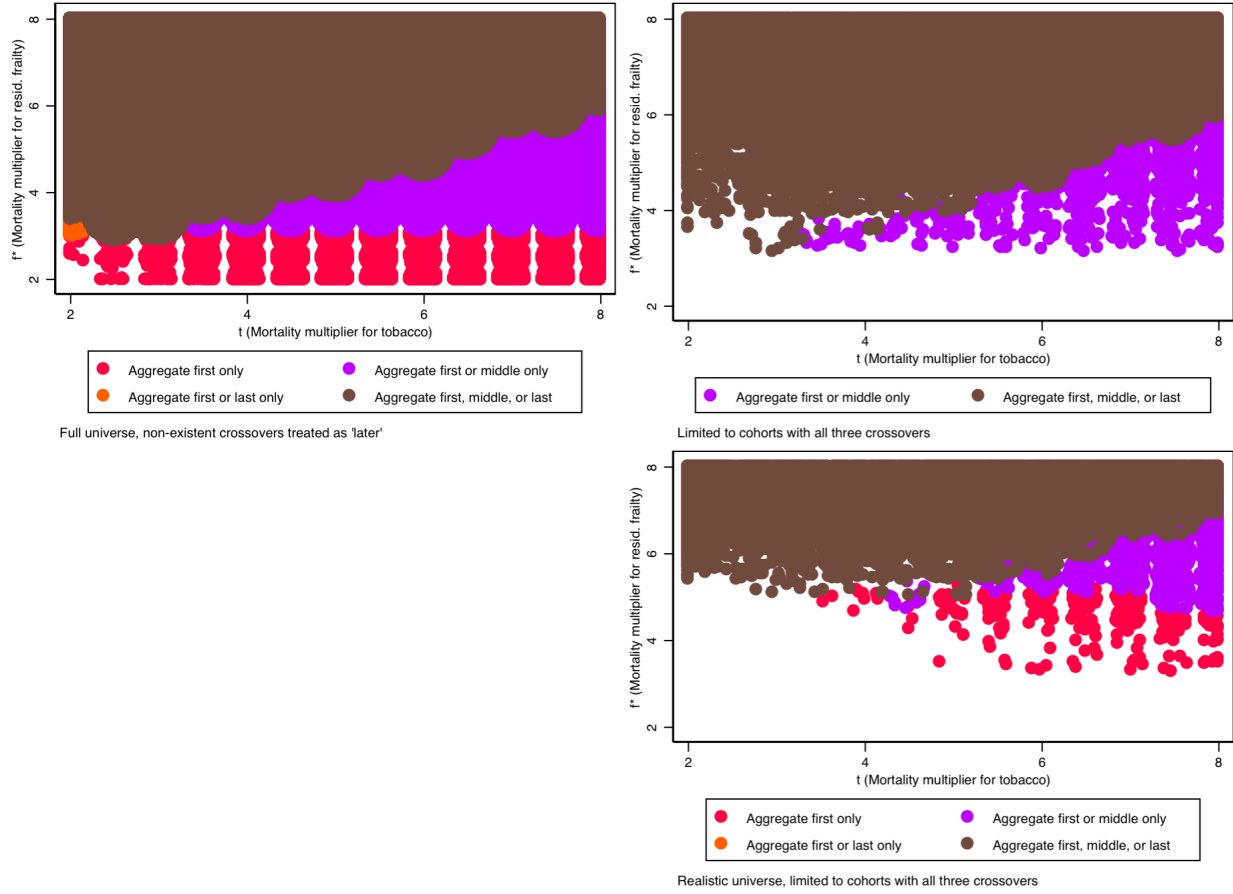


Figure S3.1 Crossover order in cohort sets defined by mortality multipliers. The cohort sets shown here all share a black mortality multiplier of $b = 2$.

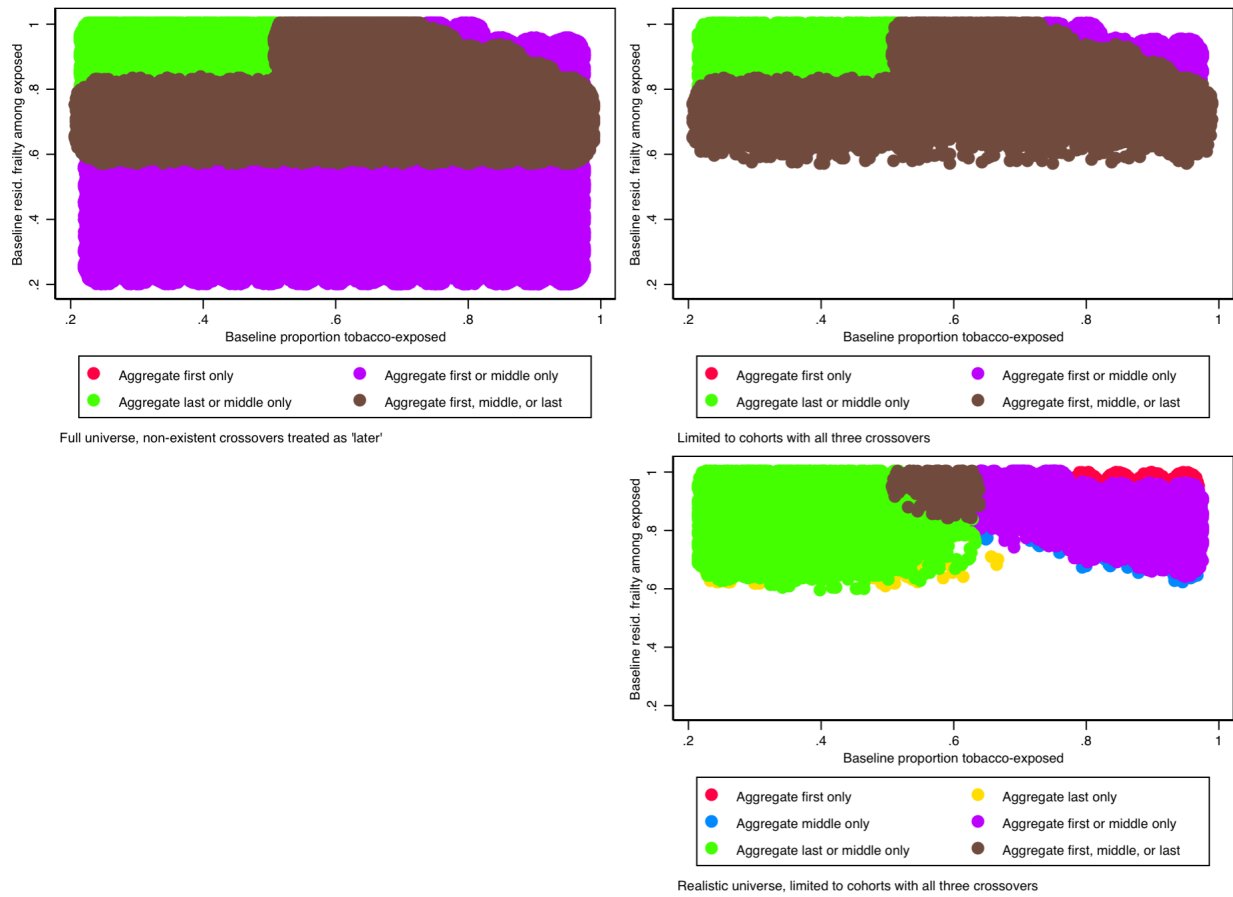


Figure S3.2. Crossover orders in cohort sets defined by baseline heterogeneity distributions. The cohort sets shown here all share a baseline proportion frail among the non-exposed of $\pi_n(0) = .75$.

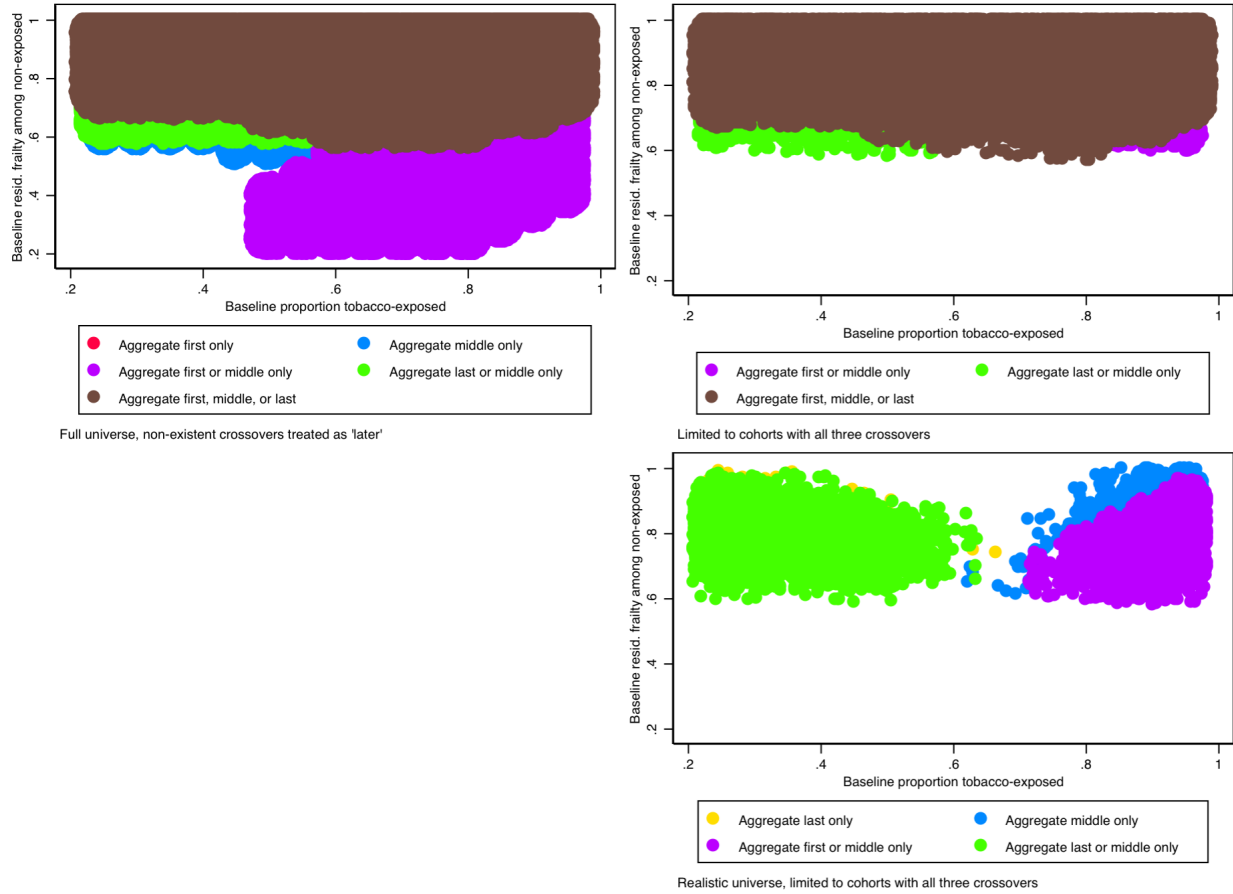


Figure S3.3. Crossover orders in cohort sets defined by baseline heterogeneity distributions. The cohort sets shown here all share a baseline proportion frail among the exposed of $\pi_t(0) = .75$.