A multi-stage extension of mortality selection theory to explain mortality deceleration among U.S. elderly subpopulations

Elizabeth Wrigley-Field and Felix Elwert
University of Wisconsin at Madison
December 31, 2009 Draft

Contrary to what was long demography's conventional wisdom, many populations' mortality rises more slowly at very advanced ages than the exponential increase it displays throughout most of adulthood. Such *mortality deceleration* has been documented in human populations (Olshansky 1998; Vaupel et al 1998; Horiuchi and Wilmoth 1998; Lynch and Brown 2001; Lynch et al 2003; Kannisto et al 1994; Fukui et al 1993) as well as medflies (Carey and Vaupel 1992), nematode worms (Vaupel et al 1998), and even automobiles (Vaupel 1996).

The dominant explanation in the literature is *mortality selection*: over a cohort's lifespan, those most prone to dying do so, leaving a remaining population with a correspondingly diminished average hazard. In the language of mortality selection, the *frailest* die, leaving a more *robust* population at old ages.

Traditional mortality selection theory implies that high mortality populations should decelerate at younger ages than their lower-mortality counterparts, because the former experience more intense selective pressure. Indeed, a major theoretical result in the mortality selection literature is that even when the frail and robust subpopulations' mortality are proportional in two populations, the one with greater absolute mortality will be more heavily selected for robustness (Vaupel and Yashin 1985; Vaupel, Manton and Stallard 1979; Kannisto 1994). This result has been used, for example, to explain the black-white mortality crossover in

the United States: even when the two populations are assumed to have equal baseline frailty distributions, African-Americans' higher mortality throughout most of the lifecourse can make them so much more selected than white Americans that at the oldest ages, their average mortality is lower than whites' (Manton and Stallard 1981; Lynch and Brown 2001; Dupre 2006; Elwert and Wrigley-Field, unpublished).

The conventional expectation that higher-mortality populations will decelerate at earlier ages has been borne out empirically – for example, in the finding that the age at onset of deceleration has risen over time as cohort mortality rates have fallen (Horiuchi and Wilmoth 1998; Lynch et al 2003) – with one exception. Lynch et al (2003) find that African-American mortality decelerates at older ages than white Americans'. They explain this counter-intuitive finding with the hypothesis that African-American cohorts have a larger proportion of frail members from birth, so that even if they are more heavily selected, they also must accumulate more selection to decelerate significantly.

This paper reconsiders whether mortality selection can plausibly explain differences in the timing and extent of mortality deceleration among U.S. subpopulations. Using high-quality, near-population-level U.S. Medicare data and nearly nonparametric hazard models, we begin by reexamining racial and sex differences in deceleration. Unlike Lynch et al (2003), and in line with traditional mortality selection theory, we find that deceleration begins earlier and extends farther among the higher-mortality populations: African-Americans and men. Indeed, we find some evidence of a *decline* in the mortality hazards of black men, beginning in the mid-90s ages.

When we extend this analysis to important acquired dimensions of heterogeneity in mortality experience, however, we get a very different result. We examine the relative age and extent of deceleration among populations stratified by baseline health and baseline poverty

status. To our knowledge, this is the first empirical examination of deceleration along dimensions of socioeconomic status and health. We find that, with the exception of white women, it is among the non-sick, non-poor that mortality decelerates substantially earlier and farther than in their higher-mortality counterparts.

While this finding is at odds with the standard predictions of mortality selection theory, we extend selection theory to account for it. Our model of *multi-stage mortality selection* shows that the higher-mortality subpopulation is not always the one facing the greatest selective pressure – if frail individuals are more likely, not only to die, but to switch from one subpopulation to another. These empirical and theoretical results suggest the need for greater substantive knowledge about inequality and the processes of acquiring health disadvantage, to make more precise predictions about what patterns of mortality deceleration might be generated by mortality selection, and ultimately, to evaluate whether selection can successfully explain otherwise-counterintuitive mortality patterns.

DATA & METHODS

Data

We analyze a large, longitudinal dataset derived from Medicare Claims Databases from 1993 to 2002.[1] We follow 28.7 million Americans over those nine years. The major benefits of these data are their near-population coverage, their accuracy, their precision, and their inclusion of covariates representing important dimensions of frailty.

Medicare databases capture 96 percent of Americans above age 65 in 1993. We restrict the analysis to individuals aged 70-97, covering the age of deceleration while avoiding sparsely populated extreme ages and cohorts with problematic age reporting (Owens and Parnell 1999).

[1] We thank Nicholas Christakis and Laurie Meneades for assistance with data development and management.

The accuracy of age reporting in the data stems from the Social Security Administration (SSA)'s rigorous monitoring to prevent fraudulent benefit claims. Such accuracy is crucial to any evaluation of mortality deceleration, since overestimation of the oldest ages may have marred some attempts to measure deceleration (Elo and Preston 1994; Preston et al 1996, 1999; Kannisto 1994). In a comparison of an earlier period of Medicare, Census, and death certificate data, Kestenbaum (1992) finds the Medicare data to be the most accurate.

The ages, moreover, are precise, with exact birthdates and daily death date followup (Sohn et al 2006, Elwert 2008). These data, drawn primarily from the SSA's Master Benificiary Record file, appear to surpass all other national mortality datasets -- even the SSA Death Master File -- in matching the National Death Index (Sohn et al 2006).

Beyond age, the data include respondents' sex, race, baseline poverty, baseline health, and region of the country. Sex and race information come from the Medicare Vital Status file. The race variable is drawn from the SSA's Master Beneficiary Record and updated from self-reported race from applications for replacement Social Security cards. We limit our racial comparison to African-Americans and white Americans because previous research has supported the accuracy of those classifications (Lauderdale and Goldberg 1996; Arday et al 2000; Elwert and Christakis 2006).

Baseline poverty status is measured as joint eligibility for Medicaid and Medicare in 1992. We construct the baseline health measure to summarize detailed health information into a standardized unidimensional measure. The Medicare Provider and Analysis Review file provides in-patient hospitalization records for 1992, from which we extract physician-provided information about chronic illnesses. We summarize those in Charlson Comorbidity Scores, a weighted count of serious chronic conditions, widely used in medical research and considered a

reliable predictor of succumbing to further health stresses. Work done by Yashin (2007a, 2007b) on the Cumulative Index also suggests that morbidity measures summarizing the number of distinct health detriments, as in our Charlson scores, may be excellent measures of aging and health deterioration processes.

Since our dataset combines information from several different Medicare files, the files were matched using unique individual-level identifiers. All files were successfully matched for all individuals.

Methods

We calculate nearly non-parametric hazard rates using exposure-adjusted poisson analysis on a high-dimensional contingency table. We calculate rates that are independent for each race and sex, and enter each three-month age unit as a separate dummy variable so that the shape of the hazards over the agespan is not constrained to any particular functional form. We use three-month age units in our analysis to ensure maximum flexibility, since, particularly at the oldest ages, hazards may change too quickly to be well captured by yearly age-specific hazards (Gavrilov and Gavrilova 2009); however, we multiply the calculated rates by four to report them in the standard yearly scale.

We interact our poverty and sickness variables by each age dummy, separately for each race and sex, so that their effects also may vary freely over the agespan. Poverty is a single binary variable, while sickness is two dummy variables (corresponding to a Charlson score of 1, with the substantive meaning of moderate sickness, or a Charlson score of 2+, meaning more serious illness, both compared to a baseline of zero, meaning a lack of chronic conditions). We

also interact poverty with both of the two Charlson variables, so that, for example, moderate sickness can have a different effect for poor white women than it does for non-poor black men.

Finally, we adjust all rates for region of the country (entered as nine dummy variables) and for cohort. Thus, the respects in which our rates are parametrically constrained are the assumptions that the logged effect of each region, and of cohort, are linear over the agespan and constant over poverty and sickness (though not race and sex), and that the interaction between poverty and sickness are constant over the agespan. Of these, we considered the linear cohort effect to be the most problematic, and initially assumed it was too restrictive; to our surprise, however, the resulting rates were extremely similar, for all race, sex, sickness, and poverty groups, to rates calculated with a cohort effect allowed to vary freely over the agespan (not reported). This may reflect the relatively small number of cohorts observed in our data, nine yearly cohorts for each age; we would expect a more complicated cohort pattern over a longer stretch of time.

The presumption in all studies measuring a single point of onset of deceleration is that the hazard accelerates continually, as in a Gompertz exponential curve, until reaching its maximum and decelerates thereafter. We measure the onset of deceleration as the point of maximum acceleration, that is, the inflection point of the hazard's second derivative, and more generally examine deceleration by examining patterns in the second derivative. Rau et al (2009) argue that this is the most physically natural interpretation of deceleration, and show empirically that, among alternative measures used in the literature (namely the maximum Lifetable Aging Rate and the maximum first derivative), it is also closest to what most demographers probably mean intuitively by deceleration: the point when the hazard begins to deviate significantly from a Gompertz curve. We also, however, go beyond examining point estimates of when deceleration

begins to compare related patterns in populations' hazards, such as when their second derivative becomes negative (which we might think of as 'objective deceleration,' compared to the deceleration relative to the previous rate of acceleration that is measured by when the second derivative is at its maximum). We measure decline in the hazards with the physically intuitive definition of negative slope (negative first derivative). Since our hazards are nonparametric, first and second derivatives over age are calculated directly from the age-specific rates using Stata (separately for each population being compared), rather than from parameters of the hazard function in a parametric model.

The sheer quantity of observations in our data is a precondition for the lack of assumed parametric form, but so is the quality of age reporting in the data. Lynch et al (2003) argue that the parametric assumptions of their rates constitute a correction against misreported ages, since rates at the oldest ages are partially extrapolated from death rates earlier; they also show that their arctangent form is very similar to the standardly used logistic curve, so this argument applies generally to the parametric forms used in the deceleration literature. Yet the precise nature of the correspondence between hazards at the oldest ages and at somewhat younger ones is exactly what is unknown, necessitating empirical investigation of deceleration patterns. Thus, even relatively flexible parametric curves run the danger of understating the divergence of oldest-age hazards from their earlier trajectory by interpreting the true late-life rates as being driven by noise rather than a genuine change in the age pattern of mortality.

Moreover, the maximum likelihood method used to estimate the parameters of the parametric models used in most mortality studies will draw more heavily from early ages for some populations than for others: where there are many more observations at younger ages than at the oldest ones, the estimation of the oldest ages' hazards will be more weighted by the

younger ages than when the number of observations is more constant across age. And this, in turn, is a function of each population's hazard at the relevant ages (as well as of differences in cohort size). Thus, the estimation procedure most commonly used in studies of relative deceleration will be more biased against finding such deceleration in some populations than others. Clearly, this is a problem for comparing deceleration patterns across populations.

This problem is circumvented with our nonparametric rates, since each age's rates are calculated (nearly) independently of the others. We do, however, need to smooth the hazards, since our nonparametric estimation means that even very small fluctuations in the hazard can alter slopes and accelerations substantially. We use a lowess smoother with a bandwidth of .2. Importantly, since we smooth on a collapsed form of our dataset in which the death rate at each age is a single observation, our smoothing does not draw more heavily from early ages in some populations than others due to different age distributions. The smoothing is sufficient to remove erratic changes in the sign of the first and second derivatives while changing the hazard very little in most cases, since the sheer amount of data makes the rates unusually smooth even with narrow three-month age units. We calculate first and second derivatives from the smoothed rates, but display estimated and smoothed rates together for each subpopulation (below) so they can be visually compared.

RESULTS

Race and sex comparison

We first compare deceleration for each race and sex, aggregated over poverty and health. The estimated and smoothed mortality hazards are displayed in Figure 1, and the slopes of the hazards in Figure 2.

We begin by calculating the age at maximum acceleration (maximum second derivative) for each population. This is a measure of *relative deceleration* since it tells us when each population's mortality begins accelerating at a slower pace than previously, thereby deviating from the Gompertz curve's continually increasing acceleration. We find that men of both races decelerate before women, and that race makes surprisingly little difference:

Age at maximum acceleration (relative deceleration)	**White**	**Black**
Women	90	90
Men	88	87

Table 1. Age at maximum second derivative (smoothed hazards) for each race/sex population.

The signs of the first and second derivatives, by revealing the pattern of *absolute* deceleration, yield further insight into the population differences. White women's acceleration diminishes in magnitude during the 90s, but stays positive at all ages: in absolute terms, white women's mortality continues to accelerate at all observed ages, even though it accelerates more slowly after age 90. Black women and white men, by contrast, begin to experience *negative* acceleration at ages 92 and 93, respectively.

Black men show the most dramatic pattern: while their second derivative becomes negative at age 88, their *slope* becomes negative -- i.e., their mortality declines -- after age 95. This apparent decline in black men's mortality is an important finding because previously reported declines in human mortality have been met with skepticism about the quality of the data (see, e.g., Kannisto 1991). Studies with more carefully verified ages generally have not found hazard declines in human populations, although those studies have been based on samples -- often with small sample sizes at very old ages -- and have used parametric forms that may be able to capture mortality plateaus but not declines. Here, although we observe nearly the entire population of elderly African-American men, the population size is small enough that their

mortality rates at the oldest ages are somewhat erratic, as can be seen in their unsmoothed hazard curve. Given this, and given that this mortality decline is a unique finding among studies with apparently accurate age reporting, we cautiously conclude that African-American men's mortality decelerates at earlier ages than the other populations, and *may* in fact decline.

Table 2 summarizes the qualitative pattern of the slopes:

Age of absolute deceleration and decline in mortality	**White**	**Black**
Women	None by age 97	Negative acceleration by age 92; slope stays positive through age 97
Men	Negative acceleration by age 93; slope stays positive through age 97	Negative acceleration by age 88; negative slope by age 95

Table 2. Age of absolute deceleration and decline, as respective measured by when (if ever) the second and first derivates become negative.

These results accord with the prediction of conventional mortality selection theory: the population that experiences the highest mortality during most of the lifecourse – black men – decelerates the earliest and most, while the population with the lowest mortality at most ages – white women – decelerates last and least.

Baseline health and poverty comparisons

Next we conduct the same analysis on subpopulations defined by their baseline health and baseline poverty status, as well as race and sex. The mortality hazards (estimated and smoothed) for each population are displayed in Figures 3-6, the hazards' slopes in Figures 7-10.

Our key finding, readily seen in the graphs, is that – with the exception of white women – deceleration is earliest and sharpest among the non-poor, non-sick. For men, not being poor is the most important predictor of early deceleration, with even sick men who are not poor decelerating at earlier ages than poor men of either race and sickness status. Among women,

sickness is most important, though the pattern reverses across race: it is the non-sick black women and the sick white women who decelerate earliest.

		White		Black			
		Non-poor	Poor	Non-poor	Poor		
Women	Non-sick	97	97	**89**	**94**	Non-sick	Women
	Sick	**94**	**94**	97	96	Sick	
Men	Non-sick	**90**	96	**86**	97	Non-sick	Men
	Sick	**93**	96	**94**	97	Sick	
		Non-poor	Poor	Non-poor	Poor		
		White		Black			

Table 3. Age at maximum acceleration (i.e., age at onset of relative deceleration) for each subpopulation defined by race, sex, poverty, and Charlson score. The lowest ages for each race-sex population are bolded, showing that they occur among the non-poor for black and white men, the non-sick for black women, and the sick for white women.

Further, negative slopes (mortality decline) occur for only two subpopulations: non-sick, non-poor black men from the ages of just below 95 to 97 (nine consecutive three-month age units), and sick, non-poor black men in the latter half of age 96 (the last two three-month units). Decline thus occurs among the least advantaged population in terms of race and sex, but the advantaged population in terms of poverty (especially for men) and sickness (especially for women).

Black men's mortality decline in the rates aggregated over poverty and sickness seems to be driven by the decline among the non-poor/non-sick in two ways: by directly contributing to lowering the average rate, and by increasing inequality between the sick/non-sick and poor/non-poor, thereby intensifying mortality selection against sickness and poverty. Here, then, is part of the explanation for black men's pattern, although why this decline for the non-sick, non-poor occurs for black men alone among the populations requires explanation.

These results contradict mortality selection theory's prediction that the highest mortality subpopulations, having experienced the greatest selective pressure, should decelerate first. Only among white women do we find that to be the case. If mortality selection is to explain mortality deceleration in full, this puzzling finding must be explained.

DISCUSSION: MULTI-STAGE MORALITY SELECTION

Following Lynch et al (2003), who hypothesized that African-American cohorts are born with a larger proportion of frail members than white cohorts to explain their surprising result, we consider whether processes besides differential mortality selection might explain an uneven distribution of frailty across socioeconomic and health groups. We argue that acquired disadvantage, such as sickness and poverty, may not adhere to the rule of thumb that higher mortality populations have faced greater selection, for two reasons.

First, becoming poor or sick is a pathway into death for a significant part of the population. The sick or poor may face greater mortality selection for robustness *once they are sick or poor*, but the non-sick, non-poor have *already* faced a form of selection for robustness: they have survived the danger of becoming sick or poor. That selection into sickness and poverty – an analogue to mortality selection – is a second route by which the frail may disproportionately exit the ranks of the non-sick, non-poor living population. This is *multi-stage mortality selection*: some people start out neither sick nor poor and are 'selected into' poverty or sickness before being 'selected into' death. If the frailty that kills people across lines of poverty and sickness also causes some of them to develop chronic illnesses or to fall into poverty, then this first stage of selection will create a sick, poor population that is frailer than the remaining population.

This theory can explain the "disadvantage-deceleration puzzle" in our empirical results: the divergence between demographers' predictions that a population with greater mortality will decelerate earlier, and the empirical observation that – while this is true along some dimensions of disadvantage – for acquired disadvantages like chronic illness and poverty, precisely the opposite seems to occur. If black men are far more selected into sickness and poverty than white women, then this will be one of the forces elevating black men's aggregate mortality above white women's, increasing the selective pressure exerted by mortality on the whole population of black men. Yet those same elevated rates of acquiring sickness or poverty will produce earlier deceleration among the non-sick/poor, as that becomes an ever-smaller, ever more robust portion of the elderly black male population. And, as we showed with black men, that will also tend to increase inequality between the sick and non-sick (and poor and non-poor), accelerating mortality selection's tilting of the population toward the non-sick/poor. Multistage mortality selection can thus explain why the fixed-trait population with greater disadvantage earlier in life (e.g., black men) will see younger deceleration, and why that deceleration will occur among its subpopulations that are defined by being robust enough to have avoided acquiring additional disadvantages.

A priori, there is no reason to suppose that the selection for robustness of the non-sick/poor (who have avoided either dying or becoming sick/poor) outweighs that of the sick/poor (who have avoided dying under more strenuous conditions), or vice-versa. Which subpopulation will have the greater number of frail members at a given age will depend on the details. We speculate that for white women, facing less selection into poverty and illness, the heightened mortality selection among the sick outweighs the selection of the frail into sickness, resulting in an earlier deceleration among the sick than the non-sick for that population only.

The second reason that those who have acquired disadvantage may not have survived intensified mortality selection is that they may not have been so disadvantaged for very long. The sick/poor population at any given time will include people who are newly sick/poor. The more intense are both stages of multi-stage selection – the faster the sick/poor are dying, and the faster more people are becoming sick/poor – the greater will be the proportion of newly sick/poor in relation to longstanding ones. Unlike the selection of the frail into sickness and poverty, this does not by itself create a sick/poor subpopulation that is *more* frail than the remaining subpopulation, but it does mitigate the extent to which mortality selection has made it *less* frail.

Thus, whereas we should in general expect greater selection – and thus earlier deceleration – in the higher-mortality of two groups defined by a trait that is constant at the individual level, this is not the case for acquired disadvantages. At least, it is not the case for disadvantages whose acquisition is in part a function of the same traits (the frailty) that causes death.

CONCLUSION

Contrary to a prior empirical study (Lynch et al 2003), but in line with traditional theories of mortality selection (Vaupel and Yashin 1985; Vaupel, Manton and Stallard 1979), we find that along lines of race and sex, mortality decelerates earlier and more sharply among the more disadvantaged populations. This is particularly apparent when comparing absolute deceleration – the sign of the second derivative – although it can also be seen in the ages of onset of relative deceleration, that is, the age of maximum acceleration (after which mortality accelerates more slowly). Our most dramatic finding in this stage of the analysis is a decline in the hazard of black men, although given somewhat erratic mortality hazards stemming from a relatively small

population size, as well as prior work on misreported ages, we are cautious about accepting this result. Given that such a result might have been masked by the small samples and parametric mortality hazard estimation used in prior research, however, our finding at least merits further research. At a minimum, we conclude that black men evince earlier and greater deceleration than the other populations, with white women decelerating the least (and not at all by age 97 in absolute terms).

For sickness and poverty, this pattern, for most populations, is reversed: except among white women, the non-poor, non-sick decelerate earlier than the higher-mortality sick and poor. We propose a new model of *multi-stage mortality selection* to explain these results. In this model, frail members of disadvantaged populations are disproportionately selected into illness and poverty as well as death.

Demographers have used the theory of mortality selection to explain surprising outcomes in many domains, not just mortality – ranging from declining prospects of finding a job the longer one is unemployed, to theoretical models in which marriages' "seven-year itch" is a mere artifact of population heterogeneity (Vaupel and Yashin, 1985). When the outcome is mortality, however, frailty has a simple substantive interpretation: propensity to die. Yet the same things that propel one toward death may also raise the probability of other adverse life events; indeed, such events may become the proximate causes of death. The unobserved frailty and the measured disadvantages that all tend to cause death need not be independent of one another.

This fact poses a problem for the simplest applications of mortality selection theory to populations that individuals may move between. For such populations, accurate predictions depend on the totality of ways that frailty can predispose people to leave a population – not only by mortality. Without considering each stage of selection, we may too broadly apply simple

principles ("the higher-mortality population is more selected against frailty"). Multi-stage selection, by clarifying multiple ways that the processes by which the frail move toward death alter population composition, is a useful framework for assessing the predictions of mortality selection – and ultimately, whether selection is capable of explaining the order in which populations' mortality decelerates.

FIGURES

FIGURE 1. Estimated and smoothed mortality hazards, by race and sex.

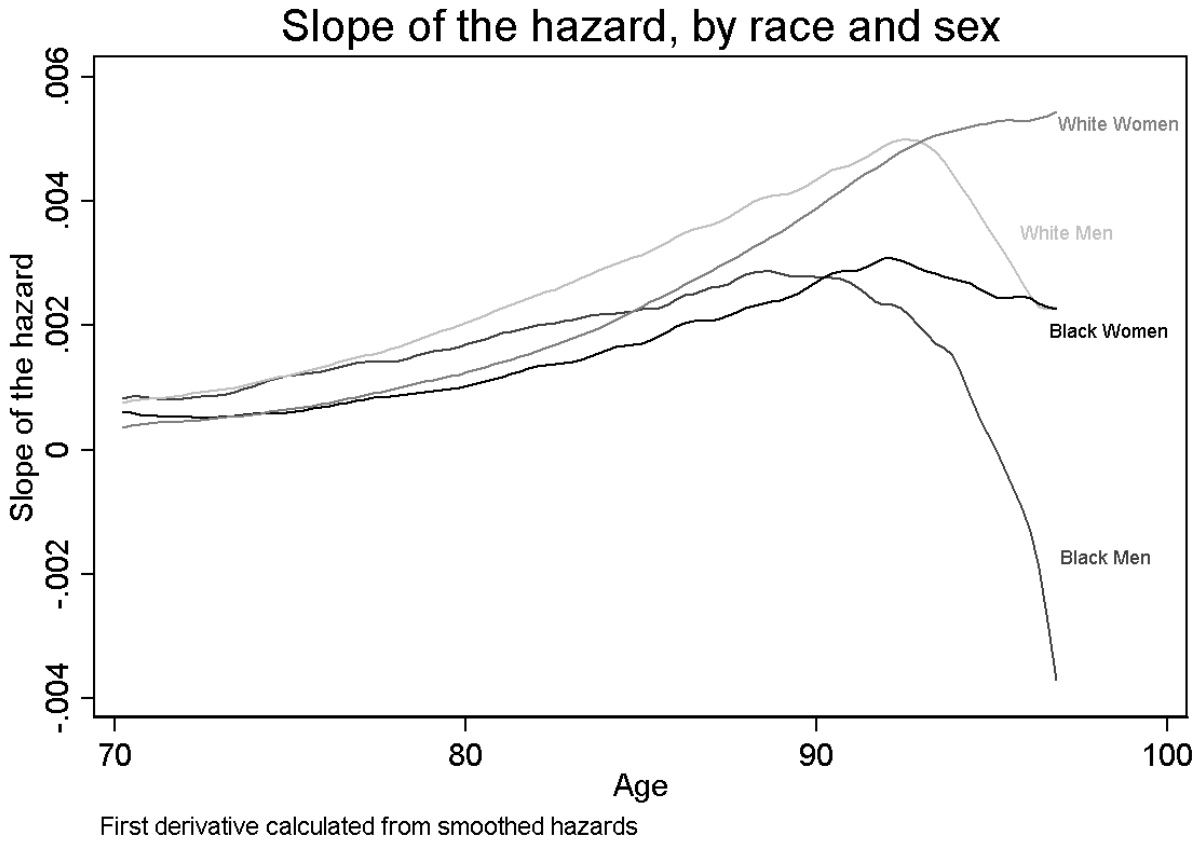

FIGURE 2. Slope of the mortality hazard, by race and sex.

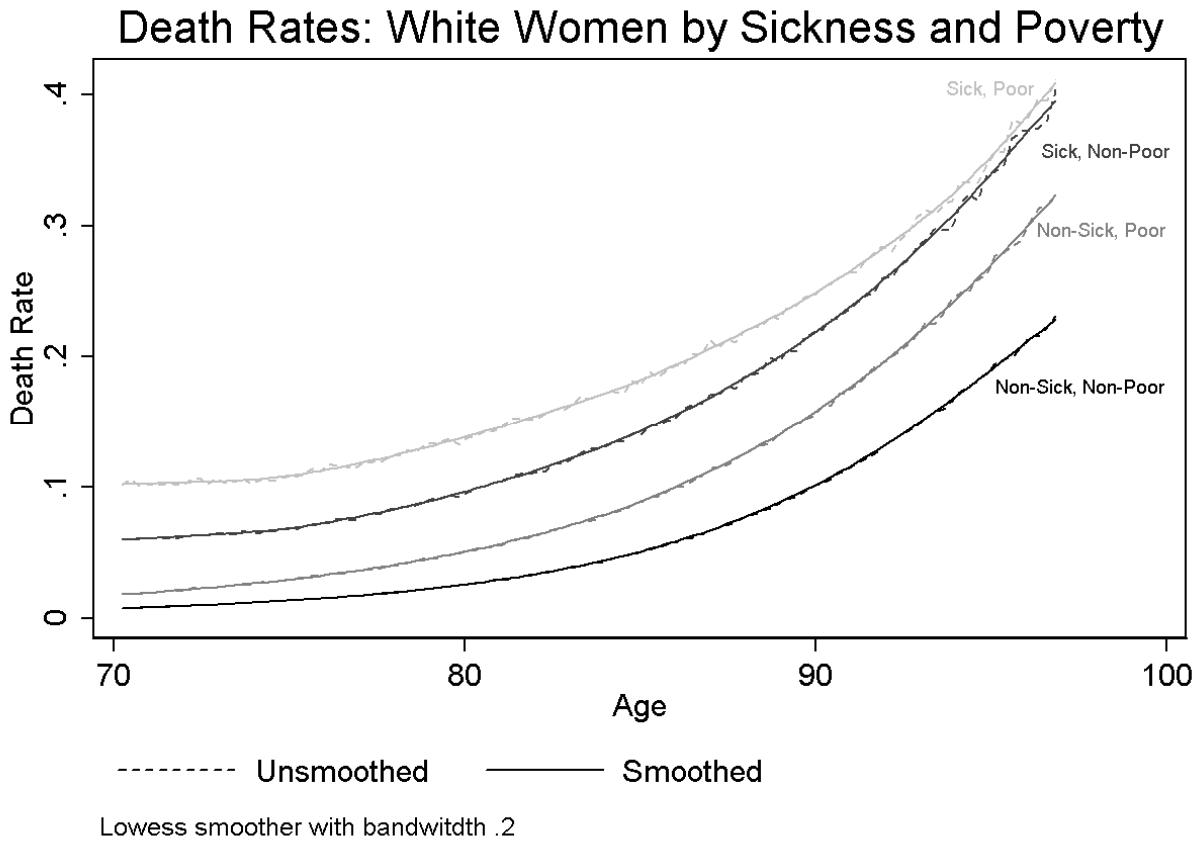

FIGURE 3. Estimated and smoothed mortality hazards for white women, by baseline health and poverty status.

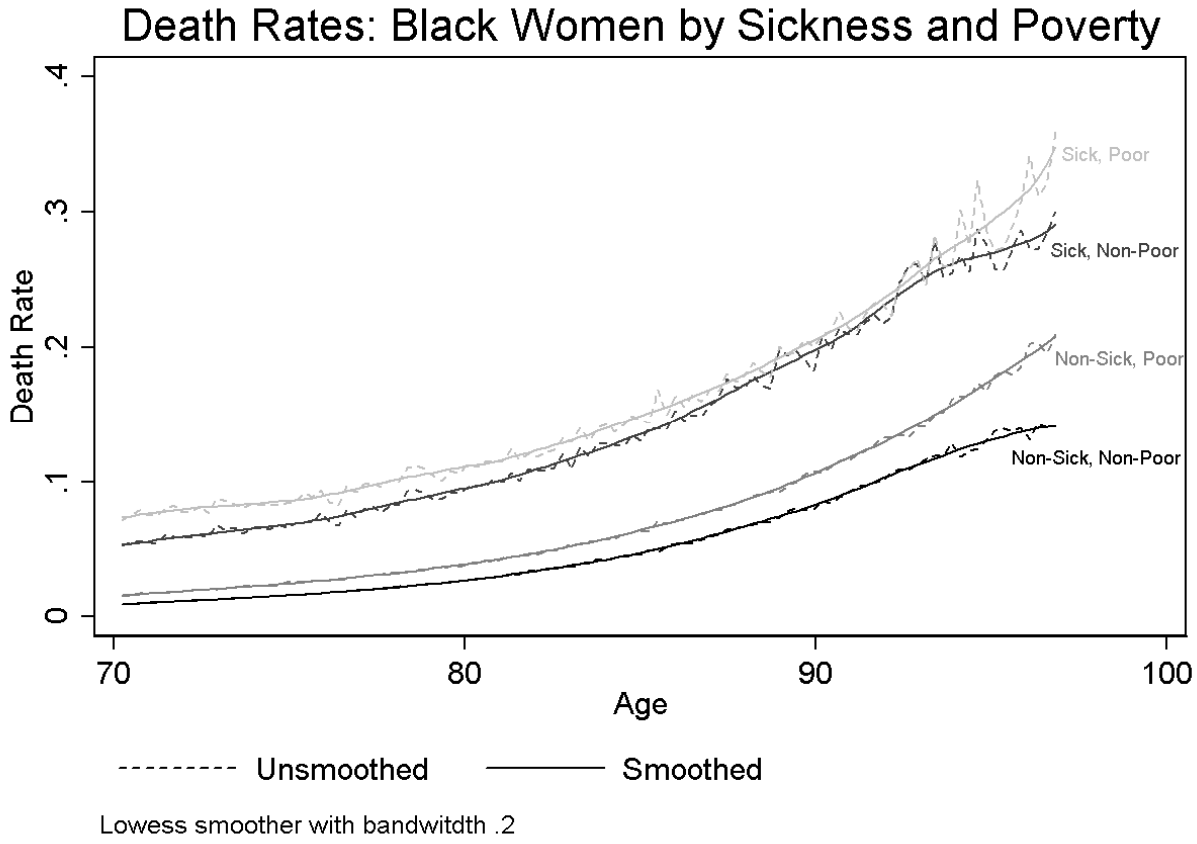

FIGURE 4. Estimated and smoothed mortality hazards for black women, by baseline health and poverty status.

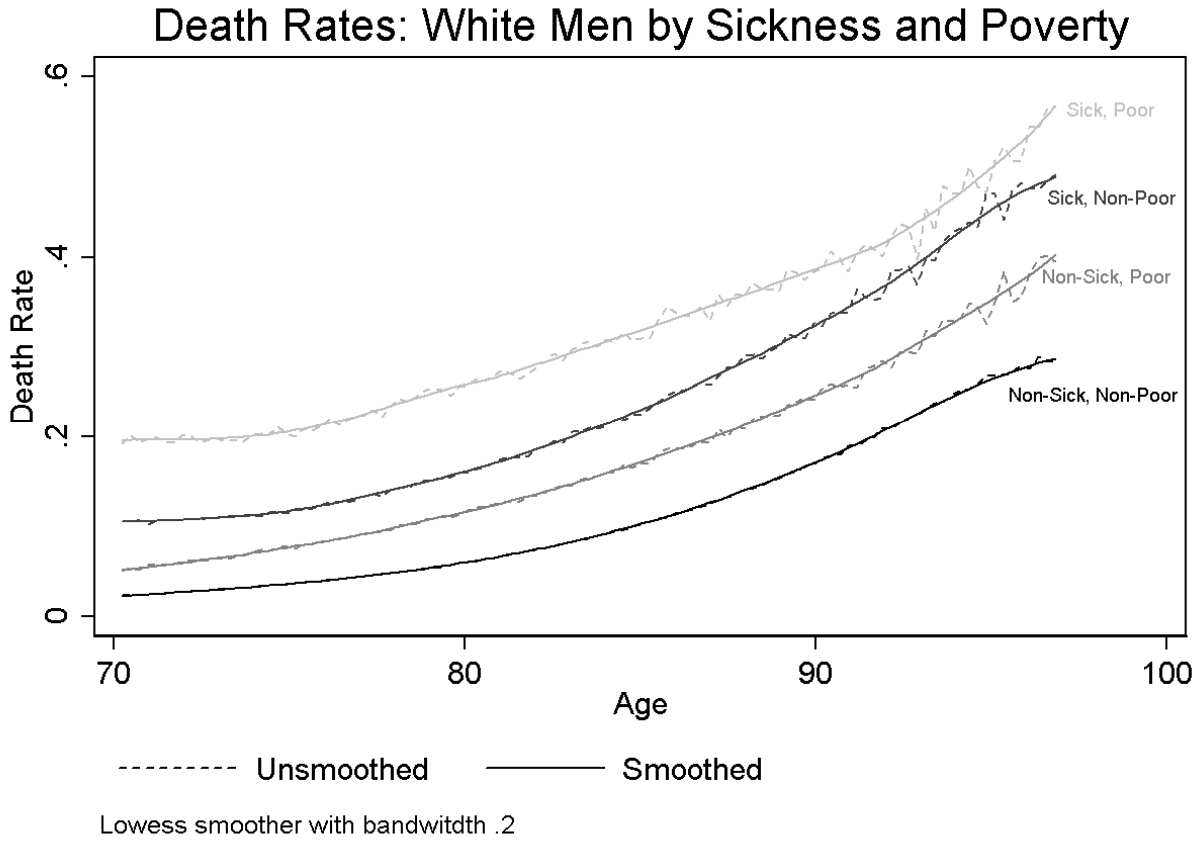

FIGURE 5. Estimated and smoothed mortality hazards for white men, by baseline health and poverty status.

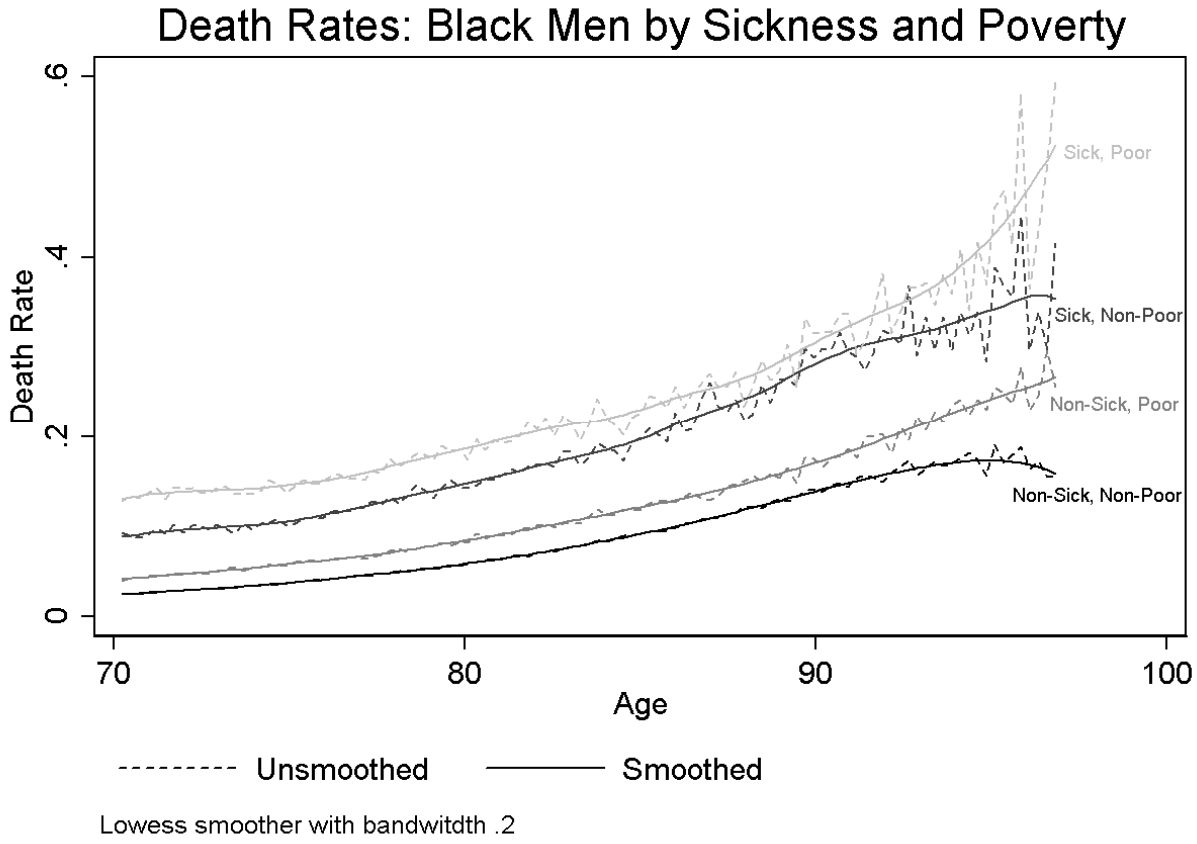

FIGURE 6. Estimated and smoothed mortality hazards for black men, by baseline health and poverty status.

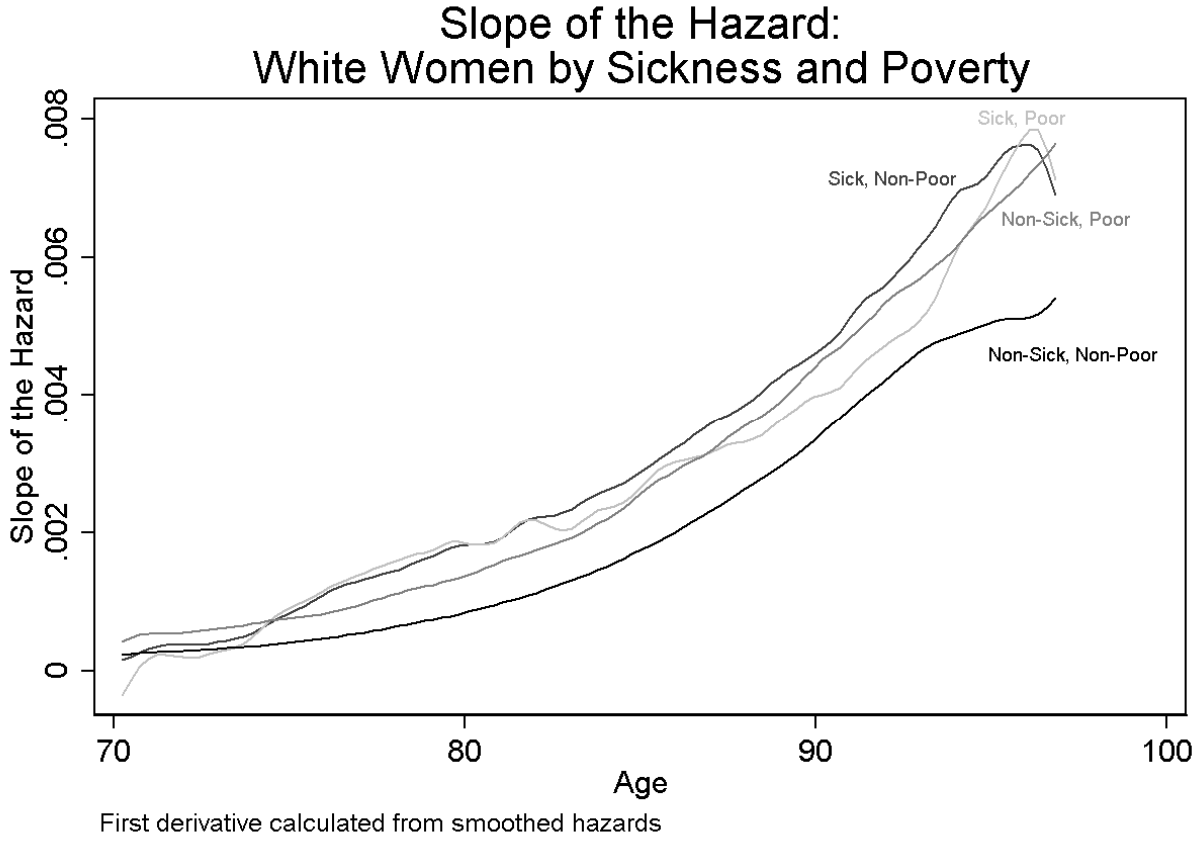

FIGURE 7. Slope of the mortality hazard for white women, by baseline health and poverty status.

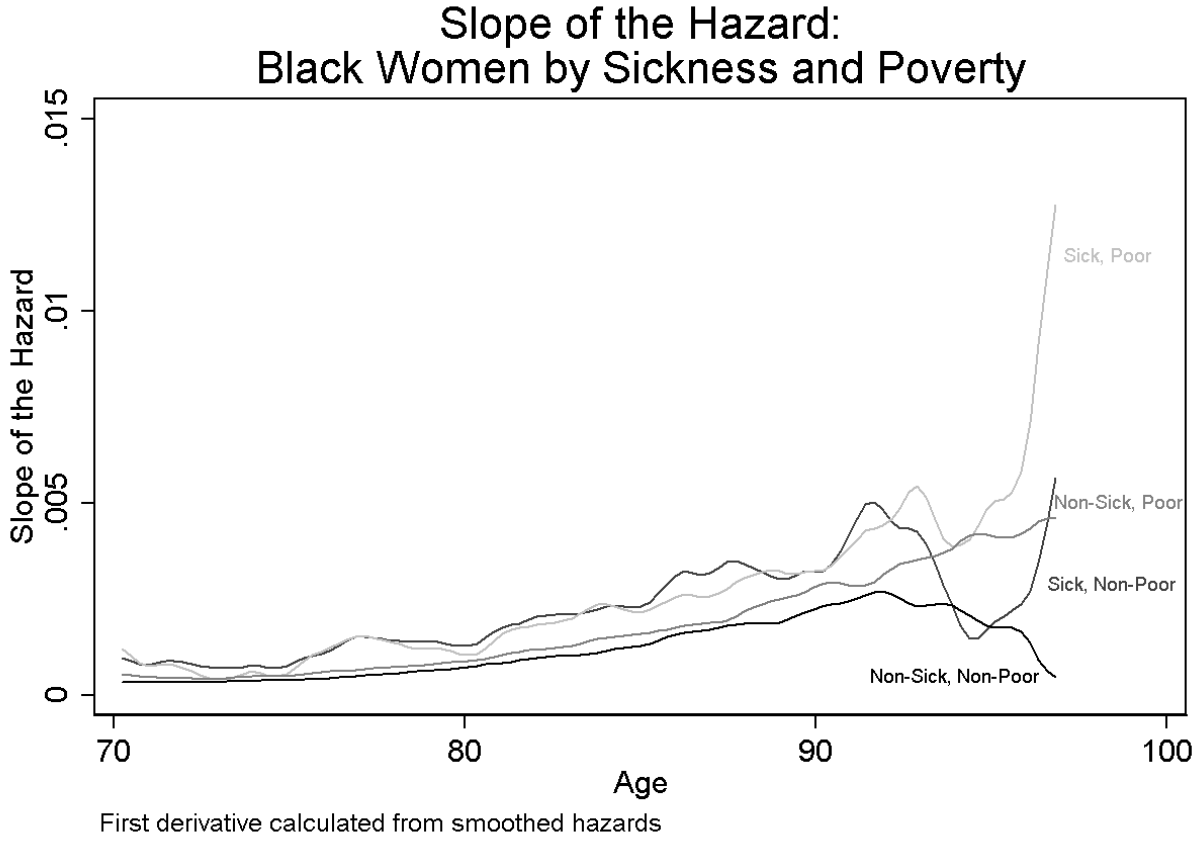

FIGURE 8. Slope of the mortality hazard for black women, by baseline health and poverty status.

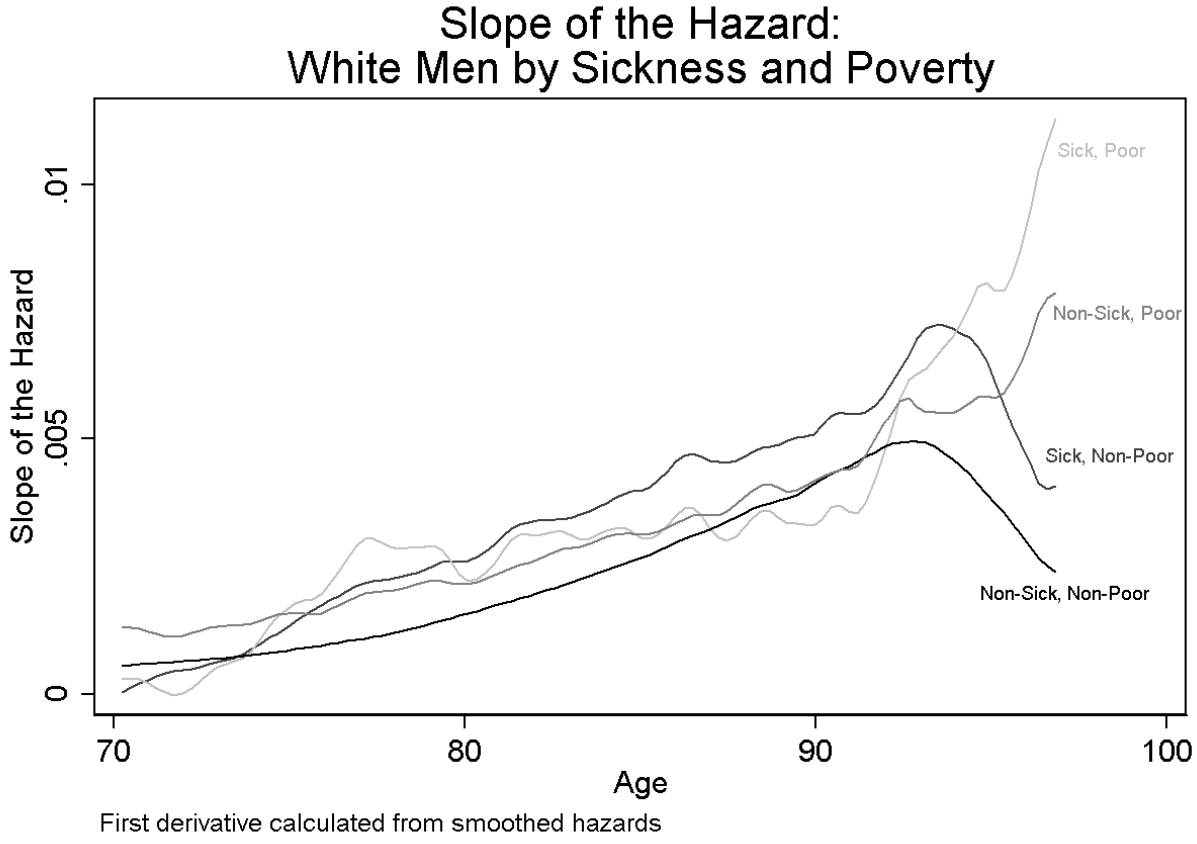

FIGURE 9. Slope of the mortality hazard for white men, by baseline health and poverty status.

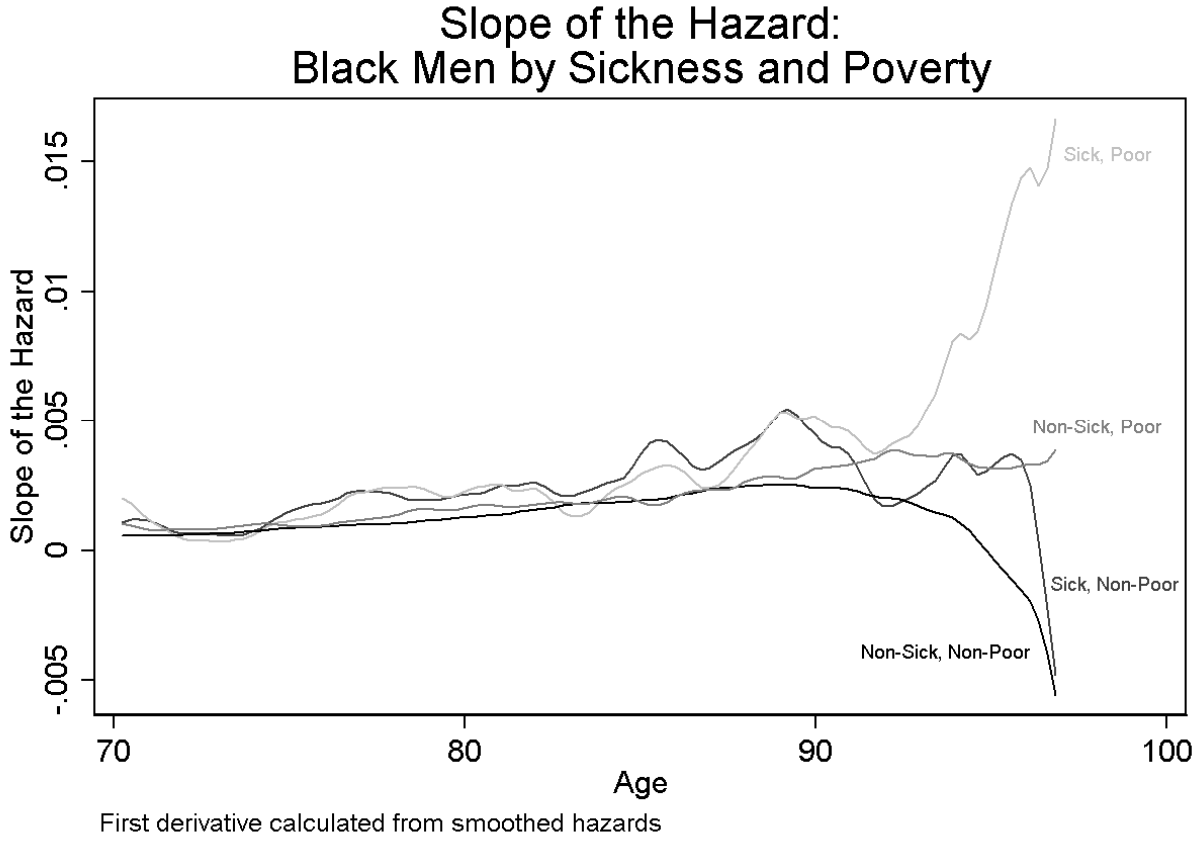

FIGURE 10. Slope of the mortality hazard for black men, by baseline health and poverty status.