A Unified Law of Mortality: Implications for the Long Run Effects of Early Conditions

Adriana Lleras-Muney and Flavien Moreau*

October 4, 2018

PRELIMINARY AND INCOMPLETE–PLEASE DO NOT CITE.

Abstract

How do social and economic conditions experienced early in life affect the evolution of health and mortality rates over the lifetime? To answer this question, we build and estimate a simple dynamic model of health. A key insight of the approach is that if mortality depends on health, then the evolution of mortality rates by age places constraints on the evolution of the underlying distribution of (unobserved) health. So mortality rates can be used to infer how health has evolved over time and across countries. We estimate our model using high quality cohort life tables from the Human Mortality Database. We use the model and the estimated parameters to understand how unexpected shocks, like wars and infectious disease epidemics experienced early in life, affect the age-profile of health and mortality. We also investigate implications for SES gradients and optimal health care expenditures.

JEL: I10, J11

Keywords: Mortality, health, evolution, life course.

*Corresponding author: alleras@econ.ucla.edu. We are very grateful to David Atkin, Andy Atkeson, David Cutler, Jeff Ely, Price Fishback and Bo Honoré for their advice and to seminar participants at Stanford University, University of British Columbia, University of Connecticut, USC, and UCLA and the NBER Cohort meetings. We are very grateful to Hualei Shang for his excellent work as a research assistant for this project. This work used computational and storage services associated with the Hoffman2 Shared Cluster provided by UCLA Institute for Digital Research and Education's Research Technology Group. All errors are our own.

1 Introduction

A large literature in medicine, demography, sociology and economics documents that circumstances early in life affect health and mortality throughout the lifetime (for recent summaries see Almond and Currie, 2011 and Almond et al., 2017). But there are several important gaps in our knowledge. For example the effects of early conditions are not always visible in the short term but appear to emerge later on in life. Sometimes the impact of policies fade and then re-appear. Thus the impact of early conditions varies depending on the age at which they are measured, but there is no well known explanation for these patterns. The extent to which selection due to selective mortality affects findings is also not clearly understood but it is often hypothesized as a potential explanation for the various findings in the empirical literature. To understand the consequences of early circumstances, and design optimal investment or compensation policies, it is necessary to have a model of how health and mortality over the lifetime are affected by inputs and insults at various ages. Yet there is no known parametric model of the production of health over the lifetime. As Almond et al. (2017) note "a structural, well calibrated model of investments in early childhood and human capital formation could help fill in the gaps in our knowledge."

In this paper we provide a unified law of mortality that tracks the evolution of health and mortality from birth to death. Cohort mortality rates exhibit a remarkably consistent pattern, high in infancy and old age, and low but variable during reproductive ages. The stability and consistency of this shape across human populations and primates, suggests there exists an underlying "law of mortality" and health (Gompertz, 1871, Carnes et al. 1996, Bronikowski et al., 2011).[1] Our model is a simple dynamic model of the evolution of the health stock that accounts for these basic features of mortality age-profiles and can be characterized by only five parameters in its simplest form. In the spirit of the classic demographic work by Vaupel et al. (1979), populations are born with an initial distribution of health (or frailty), and individuals with low levels of health die. But this distribution of health is dynamic over the lifetime. As in the seminal Grossman (1972)model, health is treated as a stock that can increase with (health) resources but otherwise deteriorates with age. But crucially, unlike Grossman's model, these resources as treated as stochastic. Finally individuals can also die from external causes or accidents, unrelated to "biological" processes and health status–this last force is not necessary to explain the basic age-profile of mortality. But accidents play an important role in explaining observed deaths during the reproductive period, while biological processes are most visible in childhood and old age. The key implication of the model is that, if mortality depends on health, then the observed age-profile of mortality imposes strong restrictions on models of the evolution of the health stock over the lifetime (among survivors). Thus the evolution of health in the population can be infered from its mortality over the lifetime.

The model makes very specific predictions about how the effects of in-utero shocks and socio-economic status affect later health and mortality and helps rationalize findings in the literature. For example the model predicts that, in the absence of compensatory responses, negative in utero shocks increase mortality at every age but the effect falls over time. However among survivors, health declines exhibit a non-monotonic pattern with age: the effects are large initially, fade by adolescence and slowly start rising with age. Similarly, permanent changes in the level (or the variance) of resources result in changes in health and mortality that vary with age in a surprising manner. Moreover these conclusions are not always the same if

[1] Most notably Gompertz (1820), Gompertz (1825), Gompertz (1862), Gompertz (1871) noted log mortality is linear after 45. Models that succesfully predict mortality from birth to death typically model the hazard rates (or some function of the rates, like survival or probabilities of dying in a given interval) using complex mathematical models. Carriere (1992) for instance shows a mixture of gompertz, Weibull, inverse Gompertz and Inverse Weibull can fit the data nicely. We provide an more in depth discussion of how our model compares to others in the literature later in the paper.

we express them in levels or in percentage changes.

We show that our model provides an excellent characterization of the age-profile of mortality for (selected) cohorts born since the early 19th century. Using the method of simulated moments and cohort life tables from the Human Mortality Database, the estimation recovers five (or more) parameters from each cohort table and can be used to predict life expectancy and conduct counterfactual simulations. We also estimate parameters for primates using data from Kohler et al. (2006) and we show that our model can also accurately describe the evolution of their mortality.

We then use the model to investigate how shocks before age 25 affect health and mortality thereafter. We demonstrate that temporary and permanent changes have very different effects on the age-profile of health and mortality. We illustrate our findings by estimating the effects of WWII, a temporary but large shock. Our model also captures well the effect of the 1918 flu pandemic, though one period shocks are more difficult to identify. Consistent with the empirical findings in the literature, we find that both WWII and the flu pandemic had long lasting "scarring" effects on mortality, despite the fact they also killed a substantial number of individuals in the short run. These effects vary depending on the age at which the negative shocks are experienced, and they are typically largest if they occur at ages with low mortality rates.

We end by investigating implications for optimal investment by age. If we assume that resources are independent of health, we find that optimal health investments are u-shaped in age: they are highest at birth, fall with age and rise again with age. Although optimal investments affect the shape of mortality, they do not fundamentally change it–mortality remains highest at young and old ages. The model also implies that lifetime investment and initial stocks are complementary, and there are also strong "dynamic" complementarities between investments at different ages, as in Cunha and Heckman (2007). We consider extensions of this optimal investment model, including making per period resources depend on health.

This paper is organized as follows. We first describe the data and the basic observations that motivate our model. We then describe the model and its properties. Then we estimate the model for several populations. We then investigate how shocks affect the evolution of health and mortality and discuss the profile of optimal investments. We finish by considering some possible applications of the model, and discuss the limitations of this work.

2 Basic mortality patterns from the Human Mortality Database

We study the evolution of health and mortality among French women born between 1816 and 1947 using data from the Human Mortality Database (HMD). The HMD provides population and death counts by age, birth-year and gender collected through vital registration systems (birth and death certificates) and censuses, from 1816 up to 2015. The availability and quality of the data for old ages is limited, so imputations are used for all ages above 90.[2] We focus specifically on French women for convenience.[3] France has the second longest time series of cohorts that can be followed from birth to age 100, after for Sweden, and it has the largest population among countries with long time-series. Using the population and death counts we compute mortality rates by age as the number of deaths divided by the population at that age (technically we are computing probability, rather than the rate, of dying at a given age). We then use these to compute

[2] These data constitute the highest quality and longest data available to study cohort mortality. But it has some important limitations. Exact population counts are only available for census years, intercensal years are estimated. Migration is not accounted for. The accuracy of the data also falls substantially for years during which the territory changed, which often correspond to wars (1861, 1869, 1914, 1920, 1939, 1943, 1945, 1946).

[3] All studies of health and mortality investigate men and women separately. A full analysis of gender differences and their evolution is beyond the scope of this paper.

the survival rates (See Appendix 8 for details). [4] For all birth cohorts we observe mortality rates from birth to age 100, except for the most recent cohorts. Cohort life expectancy was around 40 for the 1816 cohort and it rose to about 69 for the 1923 cohort, the last cohort with mortality rates up to age 90.[5]

Figure 2 shows mortality rates by age, for selected birth cohorts of French women (panel a). It plots the log (base 10) of the mortality rate by age for women born in France between 1860 and 1940. It shows that (the logarithm of) mortality has the shape of a "tick mark": it starts very high early in life, plummets to low but variable levels in adolescence and young adulthood, and then rises with age starting in middle age. If we examine the data by decade as in Appendix Figures 19 and 20 we see that this pattern is the norm across all cohorts, and is very similar (though not identical) for men. Panel b of Figure 2 shows that although there is some variation across countries, the shape of mortality is also very similar across countries for a given cohort so the evolution of cohort mortality for France is representative of the evolution in other European countries.

Several other features of the shape of mortality are noteworthy. First after middle age, log mortality rises almost linearly with age–this regularity was first noted by Gompertz in 1820 and has since led to the "search for a unified law of mortality." Second, mortality rates have declined for every age: across cohorts the curves are shifted downwards in almost parallel fashion. The steep decline in infant mortality between 1860 and 1940, could have resulted in higher mortality in older ages by leaving more frailer individuals alive.[6] But the curves do not cross–the mortality rate in old age is lower for cohorts with lower infant mortality rates, as has been noted by Finch and Crimmins (2004).

Third, the greatest deviations (in logs or proportional terms) from the tick-mark shape occurs during reproductive ages. There are visible "spikes" corresponding to war years, as can be seen for cohorts born around 1920 who experienced WWII from age 19 to age 25. Even in the absence of wars, for example for the cohorts born in 1860, there is a visible rise in mortality after age 15, which demographers refer to as a "hump" (Preston et al. 2000,Thiele, 1871),

Lastly for the most recent cohorts, there are almost no humps or visible spikes. The adoloscent hump is barely visible for the most recent 1940 cohort. The tick-shape is most clearly visible for this most recent cohort. This observation motivates our basic model which seeks to describe "natural" mortality in the absence of "external" causes that do not depend on external factors or choices such as whether to have children or events like war.

3 A parsimonious model of health and death

In this section we provide a characterization of the evolution of health and mortality based on frailty, in the spirit of Vaupel et al. (1979) (and similar to the idea of vitality in recent work by Li and Anderson, 2013). But here the distribution of frailty is dynamic over the lifetime, similar to models of in-utero shocks (Bozzoli et al., 2009 and Bruckner and Catalano, 2007). Health is treated like a stock, affected by investments and

[4] We have no information on the distribution of births and deaths within a year. So we make no adjustments for the fact that the deaths in the first year do not correspond to individuals born in the first year. While this is technically not ideal, we prefer to show the results using the fewest assumptions. The HMD reports probabilities (q_x) that make adjustments based on a series of standard assumptions in epidemiology. As Appendix Figure 26 in the shows, our naively computed probabilities are very similar to the ones they compute but they make fewer assumptions.

[5] See Appendix Figure 8. These gains resulted initially from declines in infant and child mortality–life expectancy at age 20 starts increasing substantially only towards the end of the 19th century. And life expectancy at age 40 increases substantially only after WWII. Between 1816 and 1947 France experienced several infectious disease epidemics, and three large wars (the Franco-Prussian war in 1870, WWI 1914-1919, and WWII 1939-1945), which temporarily lowered life expectancy significantly.

[6] Infant mortality in France fell from roughly 17 percent in 1860 to about 8 percent in 1940 in the HMD.

Figure 1: The evolution of mortality rates for women in France

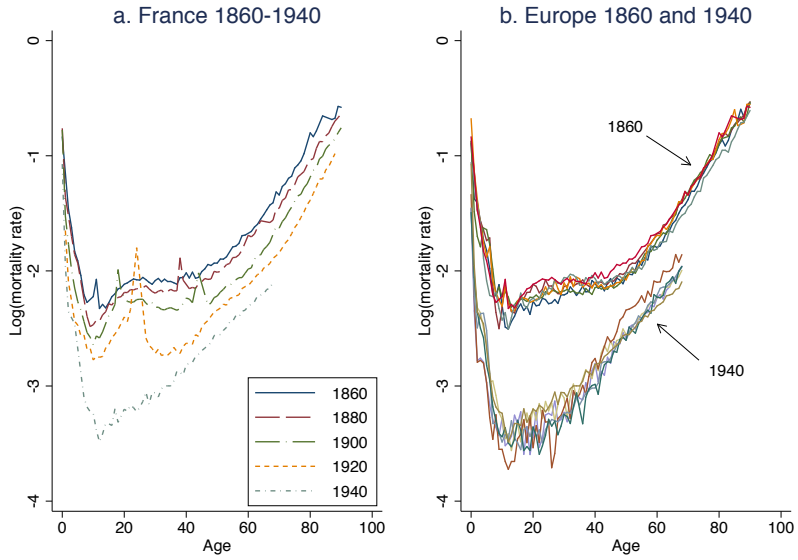

Note: Figures plot the log (base 10) of the mortality rate by age for a given birth cohort using the the Human Mortality Data. Panel a shows several cohorts of French women. Panel b shows data for women in six European countries with data for both 1860 and 1940 (Belgium, Denmark, the Netherlands, Sweden, France, and Norway). We do not show the 1816 data because only one other European country (Sweden) has data for the 1816 cohort.

subject to depreciation, as in models of human capital (Grossman 1972 and Cunha and Heckman 2007). This basic model can predict the evolution of mortality in the absence of external causes–we examine the role of external causes of death later.

3.1 A model of "natural mortality"

Assume individuals are born with an initial health level H_0. This initial health endowment differs across individuals in the population and has an unknown distribution. Every period the environment provides resources I to all individuals which increase H. In addition individuals in the environment are more or less lucky, and experience an idiosyncratic shock ε_t to their resources. For example in a stationary environment I characterizes the amount of food that a given country produces, but a given person might receive less if for instance rain was unusually low in their location. The variance of ε_t captures how unequal the distribution of resources within the population is. These idiosyncratic shocks are assumed to be i.i.d. every period. Finally the health stock is subject to depreciation every period $d(t)$ which is increasing in t ($d'(t) > 0$) : every period there is a "user cost", reflecting cumulative death cell and organ damage. Together these forces determine the evolution of the health stock.

People die when their stock of health first crosses a threshold $\underline{H}$, which is fixed throughout the lifetime and identical for all individuals. Formally let $D_t = \mathbb{I}(H_t \leq \underline{H}, D_{t-1} = 0)$ denote the random variable equal to one if the individual dies in period t. Therefore we have that the population's health and mortality can be characterized by the following dynamic system:

$$D_0 = 0$$

$$H_t = H_{t-1} - d(t) + I + \varepsilon_t \text{ if } D_{t-1} = 0$$

$$D_t = \mathbb{I}(H_t \leq \underline{H}, D_{t-1} = 0)$$

with $I \in \mathbb{R}$.[7] Note that if $D_t = 1$ then H_t is undefined–we do not observe the health of individuals if they die. Given this model, the mortality rate at time t is $MR_t = P(D_t = 1 | D_{t-s} = 0 \, \forall s < t)$.

To make this model tractable, we make a few parametric assumptions. We assume that H_0 follows a normal distribution $N(\mu_H, \sigma_H^2)$ which is consistent with the observed distribution of birth weights and other traits measured at birth (Wilcox and T RUSSELL, 1983). We will also assume that the shocks to resources every period follows a normal distribution $\varepsilon_t \sim N(0, \sigma_\varepsilon^2)$, because in simulations a normal shock provides a reasonaly fit, though it is ex-ante less clear that these are normally distributed.[8] Finally we will assume that $d(t) = \delta t^\alpha$ with $\delta \in (0, \infty), \alpha \in (0, \infty)$ which allows for the depreciation to increase over time.

In this model health is a latent **unobserved** construct that determines observed mortality. Appendix Figure 21 illustrates for the first two periods the dynamic relationship between population health and mortality rates implied by this model. The initial distribution is normal. In the first period it moves right (if I is positive and larger than the aging term) and gets wider (because of ε_t). Then the individuals to the left of the threshold, die (these individuals were either born frail or had large negative shocks). The infant mortality rate (the fraction of individuals that die in the first period) is given by the area under the curve below the threshold. In the second period this truncated distribution moves right again (if I is large relative to $d(1)$). And the population receives a new shock, generating mortality again.

Notice that in the absence of a shock there would be no deaths in period 2 – or in any period thereafter, until the depreciation term becomes large enough to push the distribution below the threshold. This illustrates that stochastic nature of the process is essential to generate mortality at every age, and it is one key feature that differentiates this model from previous ones, like the Grossman (1972) model. An implication is that the distribution of health at any age (and therefore the mortality rate) is a function of the entire history of shocks and investments. This is also clear in the definition of MR_t which conditions on survival in every previous period. The second key feature of the model is the accelerating aging component, which eventually moves the distribution closer and closer to the threshold, guaranteeing the eventual death of the entire population. This follows Grossman (1972), and is consistent with biological models of senescence (Armitage-Doll, 1954 or Pompei Wilson, 2002).[9]

Although the model is simple, it does not have closed form solutions for the mortality rate at a given age. The continuous-time analogue of our model would be a Brownian motion with a nonlinear drift, where death occurs at the first time this diffusion process hits a threshold set at zero. These kinds of models are used to model companies' default probability and to price securities in finance (e.g. Lando 2004). This literature has established that except for the particular case of a constant, linear, drift, these models do not admit a closed-form solutions for the parameters–we have an even more complex case with a potentially increasing drift. In addition we are also tracking the distribution of such Brownian motions, rather than individual ones.

[7] We could impose some restrictions on these parameters. For example the share of individuals that survive to reproductive ages is never been observed to be much below fifty percent–this would appear to be a requirement of species that do not disappear.

[8] This could be relaxed but in simulations we found that log normal errors for instance resulted in counterfactual mortality rates.

[9] See Gavrilov and Gavrilova 1991, and Weibul 1951 for attempts at biological microfoundation drawing on reliability theory from engineering.

This basic model has 7 parameters. But notice that the expression for mortality in the first period (shown in 21) is just the standard expression for the Probit model, which requires a scale and location normalization: the threshold $\underline{H}$ and the standard deviation of the initial distribution σ_H are not identified: we can subtract $\underline{H}$ and divide by σ_H on both sides of the expression determining the probability of dying and leave the mortality rate unchanged. So without loss of generality we set $\underline{H} = 0$ and $\sigma_H = 1$.[10] Thus in its simplest form, this model characterizes the biological evolution of health and mortality of a cohort using 5 (rescaled) parameters: one for the mean initial health (μ_0), two govern the aging process (δ, α), and two characterize the effects of environment, in the form of average investments (I) and the variance of these investments or shocks (σ_ε^2). We then interpret μ_H as the distance from the threshold of the initial distribution in standard deviations of the initial distribution. All other parameters are also expressed in "standard deviation" units, except for α which is "scale free"– it does not depend on the initial distribution. The proposition below shows that these normalized parameters can be identified by observing the mortality rates by age of a given cohort.

Proposition 1: The model is identified up to scale ($\sigma_H^2 = 1$) and location ($\underline{H} = 0$). (See appendix for proofs.)[11]

3.2 The behavior of health and mortality over the lifetime

We now describe the behavior of this model and then go on to analyze the effect of changes in each of its underlying parameters. Let $\hat{H}_t \equiv \mathbb{E}[H_t \mid H_t > 0]$ denote the average health in the living population with age t and $\sigma_{H_t} \equiv Var[H_t \mid H_t > 0]$ the variance of health among the living.

Proposition 2: Basic Properties of the model (See appendix for proofs.)

1. Everyone dies with probability 1: $\lim_{t\to\infty} Pr(H_t = 0) = 1$.
2. For sufficiently high I (relative to $\sigma_\varepsilon^2$and σ_H^2) mortality rates declines (up to some age t_1) and then increases with age: $MR_t - MR_{t-1} \leq 0$ if $t \leq t_1$ and $MR_t - MR_{t-1} \geq 0$ if $t > t_1$. (If I is sufficiently negative then mortality rates increase from birth onward).
3. The average health of the living increases and then decreases with age: $\hat{H}_t - \hat{H}_{t-1} \leq 0$ if $t \leq t_2$ and $\hat{H}_t - \hat{H}_{t-1} \geq 0$ if $t > t_2$ for some t_2.
4. The variance of health among the living increases and then falls: $\sigma_{\hat{H}_t} - \sigma_{\hat{H}_{t-1}} \leq 0$ if $t \leq t_3$ and $\sigma_{\hat{H}_t} - \sigma_{\hat{H}_{t-1}} \geq 0$ if $t > t_3$ for some t_3.

Figure 2 illustrates (for a specific set of parameters) the evolution of the distribution of the health stock as cohorts age in this model. This distribution at age 1 is truncated at the threshold, it moves right and broadens until age 40. Then it starts moving left and eventually becomes triangular at the threshold. At any given age after infancy and before old age, the distribution of health is very close to a normal distribution despite truncation, because it is approximately equal to a sum of normal distributions. This is consistent with the observation that health related stocks like heights, which grown from birth until maturity are

[10] More precisely we need to normalize 2 out of three parameters. We find it more intuitive to normalize the threshold rather than the initial mean, but this choice is arbitrary.

[11] We prove all propositions using the specific parametrization we used. More general proofs might be possible to derive but they would require making assumptions about the distributions. We leave this to future work.

Figure 2: The evolution of the health distribution over the lifetime

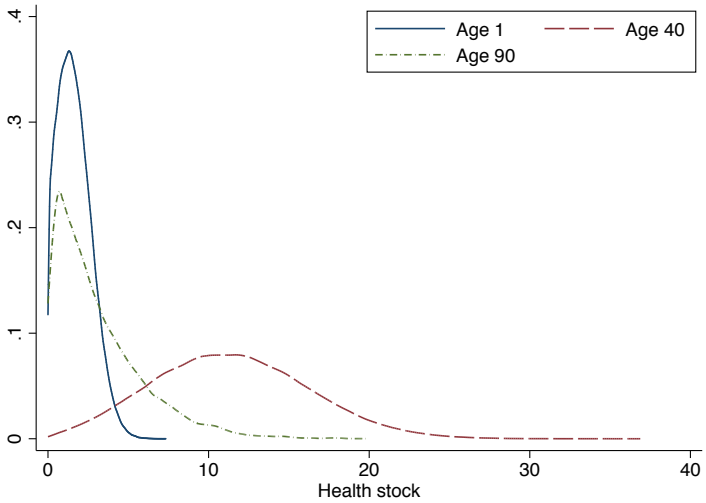

Note: Simulated data for a population of 500,000 individuals. For this simulation we use the following parameters: I=0.3575753, δ=0.0004789, σ=0.8353752, α=1.7883, μ_0=0.925079.

close to normally distributed (for example Limpert et al. 2001 show that either a normal or a log normal distribution fits female heights well).

Figure 3 shows that the model reproduces the age-profile of mortality well: (log) mortality starts high and plummets to very low levels by adolescence. It remains low and variable until around age 40, and then it starts rising almost linearly with age. The initially high infant mortality rate is mostly a result of many infants born with low health endowment, though there are also unlucky babies with large negative shocks. In childhood, mortality rates depend mostly on the the variance of the shock, and on the size of the mean investment level which pulls the distribution away from the threshold. But eventually the depreciation process becomes larger than the investment and an increasing number of individuals fall below the threshold in old age.

The figure also shows the evolution of health. Over the lifetime, health and mortality are moving in opposite directions. Average population health increases and reaches a peak in mid-life.[12] This pattern is consistent with the evolution of self-reported health by age for US cohorts (Deaton and Paxson, 1998) and UK cohorts (Contoyannis et al. 2004), and with the age-profile of productivity which depends on health. The variance of health increases and then falls, due to selective mortality. In a sense health behaves like consumption: cohort consumption inequality increases with age, so long as shocks to consumption are not perfectly correlated across individuals (Deaton and Paxson 1994, Deaton and Paxson, 1997). And then it falls because of truncation.

[12] Although we do not observe health, we can sometimes observe disease and disability rates which are also functions of health. Define morbidity as having a level of health that is above the dying threshold but below some other arbitrary threshold. In the model morbidity is also a u-shaped function of age like mortality (results not shown), being high among children, reaching low levels from ages 20 to 60, and increasing thereafter. Contemporary data also show that hospitalization days (a rough proxy for morbidity) are indeed u-shaped. For example see hospitalization rates by age for the US, which are available here https://ftp.cdc.gov/pub/Health_Statistics/NCHS/NHIS/SHS/2014_SHS_Table_P-10.pdf

Figure 3: Age profile of population health and mortality

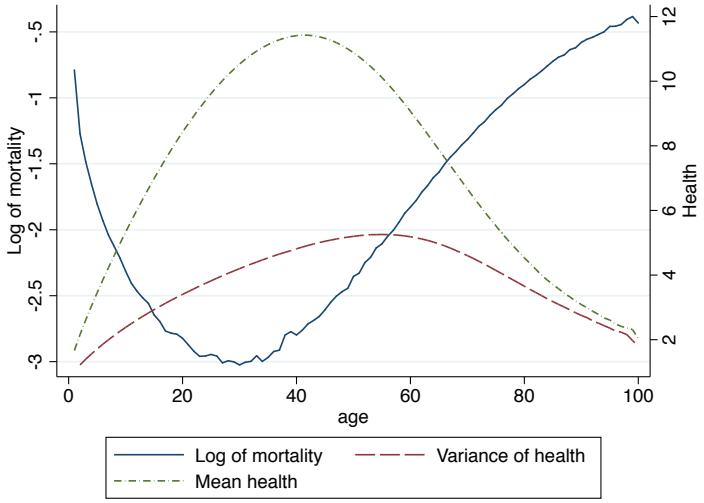

Note: Simulated data for two population of 500,000 individuals each. The baseline model is simulated using $\mu = 0.925, I = 0.35, \sigma = 0.83, \alpha = 1.78, \delta = 0.0005$.

3.3 Relationship to previous literature

There have been several attempts in the demographic literature to generate a unified model of mortality. Many models, like Gompertz's, only account for mortality after a certain age (Li and Anderson, 2013). Therefore these models do not lend themselves to a formal exploration of how early conditions affect mortality later in life. A very popular model proposed by Heligman and Pollard (1980) uses 8 parameters to describe the probability of dying at a given age for all ages. More recently Sharrow and Anderson (2016) propose a 6 parameter model, which fits the period tables well, but these represent weighted average of cohort tables. Palloni and Beltrán-Sánchez (2016) also propose a model with few parameters that tracks "Barker frailty". Our model differs in one fundamental aspect from these models: like the seminal Grossman model, we model individual's stock of health and its evolution, rather than directly modeling the mortality rate of a population. This approach allows for an easy characterization of how factors at a given age affect mortality at later age because we can model inputs into health directly.

In spirit our model differs substantively from Grossman's in two dimensions. First, as discussed above, here health evolves stochastically because of the random shocks, which are key to generate mortality at all ages. An attractive consequence relative to the Grossman model is that we do not assume a fixed horizon–the age at death is naturally determined by health and health investments in a stochastic manner. Second this model explicitly accounts for initial health conditions and traces the effect of mortality on the distribution of health among the living at any age. Thus we explicitly account for health-based selection.

Recent papers have extended the Grossman model to account for some of these features, most notably health deficit model developed by Dalgaard and Strulik (2014) and the model by Galama and Van Kippersluis (2015). Dalgaard and Strulik model aging during adulthood as a process of deterministic health deficits accumulation. This model does not match the entire lifetime mortality process, it can only match mortality starting in young adulthood and thus it is not well placed to study the consequences of early insults on later health or mortality.[13] Galama and Van Kippersluis (2015) do derive predictions about the life cycle

[13]Our model differs in other dimensions. In health deficit model, the optimal death time is the same for every individual. Instead,

Figure 4: Adding accidents to the baseline model

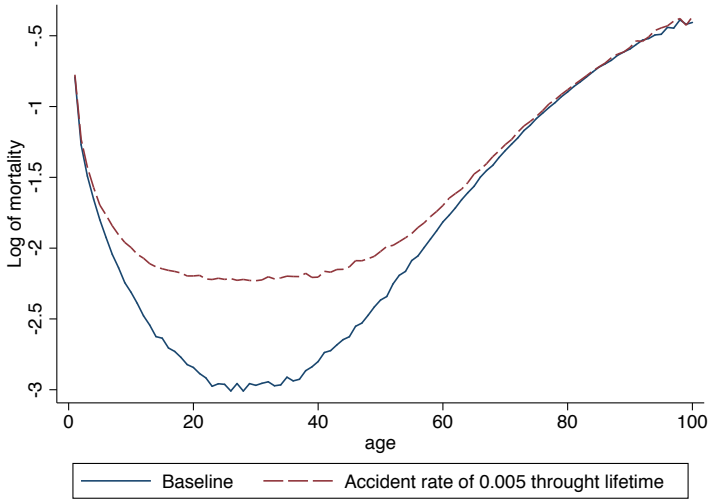

Note: The baseline parameters are the same as in Figure 3

trajectory of health by socioeconomic status, but the model is general and the predictions are ambiguous. No paper in economics that we know of has made use of the age-profile of mortality to make inferences about the evolution of health.

Our basic model is otherwise a much simpler version of the original Grossman model or these more recent papers. It does not model utility or how to think of optimal health inputs. We consider some of these issues later in the paper. But the absence of cohort data on incomes, health inputs, and their prices overtime limits our ability to empirically estimate a richer model for many cohorts. Our contribution is to note that cohort mortality life tables can be used to identify a basic parametric model of the evolution of health and mortality, upon which more complex but realistic models can be built.

3.4 Adding accidents to the model

In our baseline model, mortality is purely driven by health. However, many deaths, like accidents and homicides, strike regardless of the health status of an individual. To account for these "extrinsic" causes of death in the simplest possible way, we can extend our baseline model with an "accident shock." Suppose that a random fraction $\kappa \in [0, 1]$ of the population is killed by an accident in every period. This accident rate is assumed to be constant over the lifetime and independent of health. Each individual experiences i.i.d. shock ν_t drawn uniformly between 0 and 1 every period and dies if this shock is sufficiently high, regardless of their health.[14]

Figure 4 shows the effects of adding accidents to the baseline model. Adding accidents increases mortality rates at all ages, but more so during reproductive ages. Around reproductive ages, when biological causes of death are dampened, accidents become the dominant source of death (in percentage terms). This

in our model aging is fundamentally a stochastic process, which allows us to match closely mortality rates over the whole lifespan. In addition in our model health is hump-shaped over time in our model while it is strictly declining in the deficit accumulation framework.

[14]Intuitively this random accident rate places a floor in the level of mortality that is constant by age: if all health-related deaths were eliminated, then we would observe this accident rate at every age and its level would uniquely determine the longevity of the population ($1/\kappa$).

is consistent with empirical evidence that external causes of death (unintentional injury, suicide and homicide) account for a larger share of mortality during these ages today (Remund et al. 2018). Around birth and in old age, when the average health stock is low, accidents account for a diminishing fraction of deaths, because many individuals die due to natural causes anyway (competing risks). Although the level of mortality rises when accidents are allowed, the behavior of mortality over the lifetime is otherwise unaffected. For this reason in what follows we will start by assuming this baseline accident rate is zero.

4 The effect of permanent changes in lifetime conditions on health and mortality

We now investigate the implications of the permanent differences over the lifetime in conditions across populations. Proposition 2 summarizes the main qualitative insights, which we the illustrate graphically and discuss in light of the existing literature.

Proposition 3: Comparative statics (see Appendix 2 for proofs)[15]

1. Increasing the investment I or the average health at birth μ_H unambiguously decreases mortality at all ages: $\frac{\partial MR_t}{\partial I} \leq 0$, $\frac{\partial MR_t}{\partial \mu_H} \leq 0$.[16]

2. Increasing any of the aging parameters, δ or α, unambiguously increases mortality at all ages: $\frac{\partial MR_t}{\partial \delta} \geq 0, \frac{\partial MR_t}{\partial \alpha} \geq 0$.

3. An increase in σ_H^2 can increase or decrease the mortality rate at a given age. An increase in σ_H^2 increases the mortality rate at young ages $\frac{\partial MR_t}{\partial \sigma_H^2} \geq 0$ if $\delta t^\alpha \leq I$. Ultimately, an increase in σ_H^2 generates selection and reduces mortality in the very old age, for some t_σ, $\frac{\partial MR_{t+s}}{\partial \sigma_H^2} < 0, \forall s > t_\sigma$.

4. Investment and health at birth are complements: $\frac{\partial^2 MR_t}{\partial I \partial \mu_H} \leq 0$.

We next illustrate these effects graphically for a specific set of parameters, chosen to roughly describe the 1816 cohort. Conceptually we compare two cohorts, that are identical in all parameters except for one, which we change by 50%. Figure 5 shows how changing each parameter affects the log mortality curve. We discuss these in detail next.

Changes in any parameter affect mortality rates at all ages. Lowering initial health results in higher mortality throughout, except for the oldest (panel a) though this effect is only noticeable early on in logs (more on this below). Lowering the average annual investment also results in higher mortality at all but the oldest ages (panel b).[17] Increasing the variance (panel c) of the random shocks results in a "cross over". The population with high variance has higher mortality at younger ages but lower mortality at older ages. This occurs because when the variance is higher many more die initially. But in the population with greater variance, many individuals are also the lucky recipients of large positive shocks, and these individuals will live longer as a result. Finally increasing the depreciation rate δ results in higher level of mortality all throughout life–but the effects are imperceptible for many years, and then rise rapidly with age.[18]

[15] These statements hold in a model where the accident rate is non-zero throughout the lifetime.
[16] Changing the threshold also affects mortality rates negatively throughout the lifetime.
[17] At oldest ages there are very few individuals alive and the rates become very noisy
[18] Changing α has similar effects so we do not show them here.

Figure 5: Comparative statics for log mortality

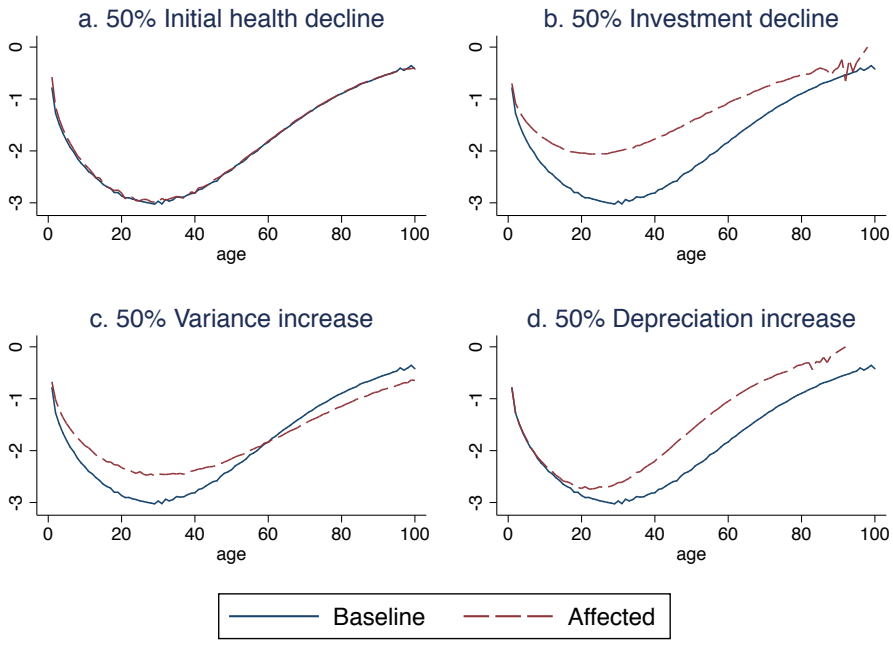

Note: Simulated data for two population of 500,000 individuals each. The baseline parameters are the same as in Figure 3

Appendix Figures 22 illustrate the effects of changing the parameters on the average health of the living. A lower initial initial stock (panel a) lowers health at all ages but particularly in old age. Lower investments (panel b) also lower the average health at all ages but the effect is not constant with age. Increasing the variance of shocks (panel c) *increases* the health of the surviving population at all ages. Lastly increasing the depreciation rate (panel d) lowers health at all ages. These results imply that cross-sectional estimates will necessarily underestimate the full effect of any of these changes, which are best summarized by how they affect (health adjusted) life expectancy.

To best match the empirical literature we present next the results as differences between affected and unaffected populations by age. The purpose is to illustrate how the effects vary with age and the extent to which conclusions match what has been found in the literature. We show results in levels and in logs (or percentage terms) for two reasons: because some papers estimate linear models and thus compute gaps in levels, while others use logs (or log odds) and report effects in percentage terms.

4.1 The effects of in-utero shocks

Negative in utero shocks such as wars, famines, disease and stress, are equivalent to lowering the mean levels of initial health in a population. Figure 6 shows the effects of decreasing initial health on mortality and health among the survivors by age. Lowering initial health results in markedly higher infant and adult mortality (except for very old ages which are not shown). The effects monotonically decline with age in percent (log) terms (panel b) but the effects are u-shaped in levels (panel a).

A lower initial stock lowers health (among the survivors) at all ages, but the pattern is u-shaped, with a large decline initially, almost no impact for many years thereafter, and increasingly larger impacts as individuals age in both levels and percentages (panels c and d). Thus the models predicts exactly what

Figure 6: Effects of decreasing initial health, by age

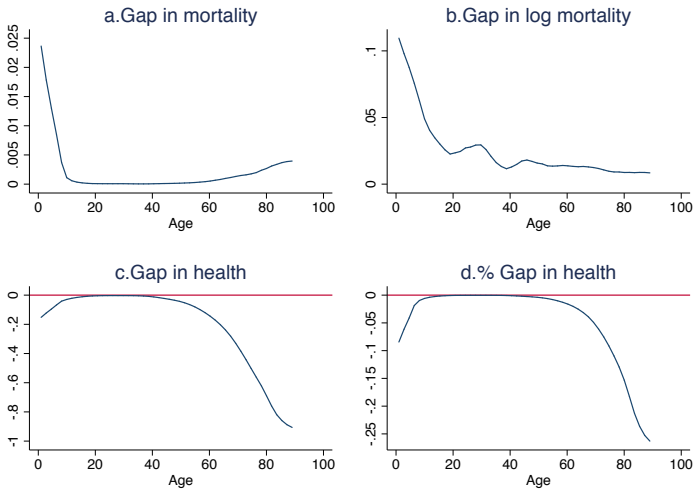

Note: Gap is computed as MR(low)-MR(high), or H(low)-H(high). The figures become very noisy after age 90 because there are almost no survivors, so we do not include these data points. Simulated data for two population of 500,000 individuals each. The baseline parameters are the same as in Figure 3

Almond et al. (2017) observe across studies: there is a "fade out" in middle age and effects re-appear later in life, and increase with age after middle age. It is also consistent with the original Barker hypothesis which predicted that events in utero would lead to chronic diseases that would only appear in later adulthood.

These predictions are in stark contrast with the predictions of the original Grossman model (in which early shocks have a decreasing effect on health with age), but in line with what studies have reported (Almond and Currie 2011). This occurs because in our model (in contrast to Grossman's) the depreciation is not a function of the level of the stock H. The model also implies that it is impossible to estimate the effects of in-utero shocks in middle age (no effects will be found), or to identify individuals that have been affected by shocks for intervention purposes during middle age. Finally note that the predictions for the age-profile of health and mortality effects are very different.

4.2 Socio-economic status and mortality.

A large literature documents large and persistent differences in health and mortality by education, income and other permanent markers of socio-economic status such as occupation and race (Cutler et al., 2012). These differences are often referred to in the literature as gradients. If education and incomes are indicative of higher average resources throughout the lifetime (I), then the model predicts that those more resources will have higher health and lower mortality throughout life except for the very oldest.

Figure 7 plots gaps in mortality and health by age between an I-rich and I-poor population. Mortality is higher for those with lower I. But the gaps are slightly u-shaped in levels: they are large at birth, decrease to almost 0 in middle age, and increase after middle age. But in log (percentage terms) the effects are hump-shaped, increasing until middle ages and declining thereafter. This results in log-mortality curves that start converging after some point (as shown panel b of Figure 6). This is very similar to the findings from Chetty et al. (2016), who show that in the US today those with high earnings at age 40 (a measure of permanent income) have lower subsequent mortality relative to those with lower incomes, with log mortality curves

Figure 7: Effects of decreasing annual investments throughout the lifetime by 50%, by age and initial health status

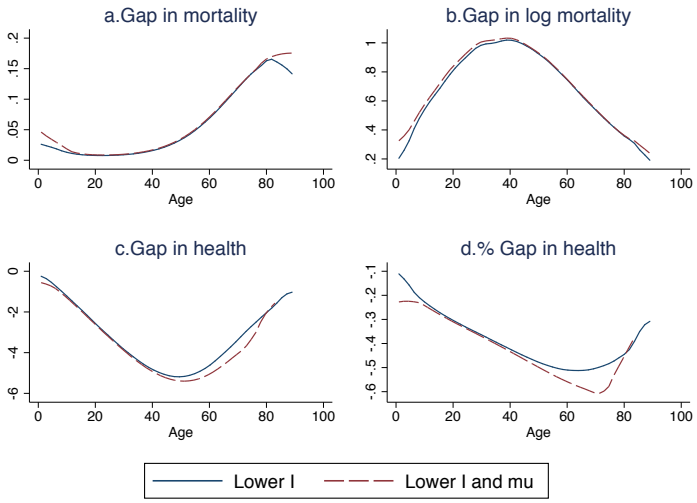

Note: Gap is computed as MR(low)-MR(high), or H(low)-H(high). The figures become very noisy after age 90 because there are almost no survivors, so we do not include these data points. Simulated data for two population of 500,000 individuals each. The baseline parameters are the same as in Figure 3

that are closer in old age. This prediction is also consistent with the literature that has documented that the effects of education on mortality fall with age, in percentage terms (Hummer and Lariscy 2011), if we view education as a strong correlate of lifetime resources. The increasing effect of I on mortality rates starting in middle age is consistent with findings in Kaestner et al. (2018) who investigate the age-profile of the effects of education on mortality (in levels).

Lower investments also lower the average health at all ages. But the effect first increases with age, and then starts declining once mortality starts rising, in levels and percentage terms. These predictions are consistent with evidence in Case et al. (2002) or Currie and Stabile (2003), who show that the gaps in self-reported health status between those born in poor and rich families grow with age, and decline after 65. Kaestner et al. (2018) also show similar evidence that education gradients in self-reported health grow between ages 30 and 65 and then appear to fall.

The figure also shows the effects of decreasing both I and initial levels of health. The combined effects are larger than the individual effects suggesting that if one does not properly control for differences in initial health, and these are correlated with levels of health resources then SES gradients or gaps will be overestimated. Interestingly the size of the bias on mortality is largest early and very late in life but small in between; but it is largest among the living at birth and after age 40.

Education and incomes could be also equated to decreases in δ, if they result in more exercise, lower exposure to pollution or lower stress, which we can conceptualize as lowering the depreciation rate. Figure 8 shows the effects of increasing the depreciation rate on mortality and health. Increasing the depreciation rate δ results in higher level of mortality all throughout life, but the effects are imperceptible for many years, and then rise rapidly with age, petering out in old ages. For health, the effects are also small before middle ages and then they rise and fall.[19] Notice that if we compare the effects of a rise in depreciation and the effects of a decrease in investment after age 40, the patterns are very similar. Thus the model predicts

[19]Changing α has similar effects so we do not show them here.

Figure 8: Increasing the lifetime depreciation rate by 50% by age

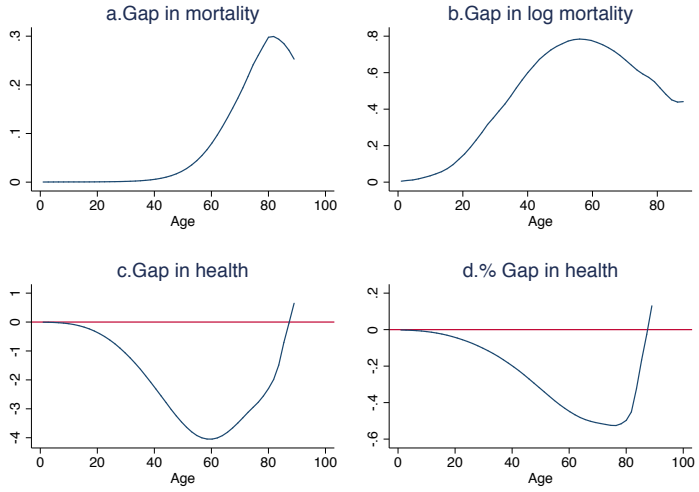

Note: Gap is computed as MR(low)-MR(high), or H(low)-H(high). The figures become very noisy after age 90 because there are almost no survivors, so we do not include these data points. Simulated data for two population of 500,000 individuals each. The baseline parameters are the same as in Figure 3

that with health and mortality data from adults *only*, it will be impossible to infer whether SES is affecting annual resources I or the annual depreciation rate δ.

The main empirical implication from this section is that estimates of the effects of socio-economic status differ depending on which parameter they affect, the age when they are measured, and the functional form that is chosen to estimate them. Moreover the bias that results from not properly controlling for initial conditions also varies with age. Last large changes in a cohort's environment (affecting intial conditions, resources, and aging) cannot always be detected until later in adulthood. Importantly, the exact age when one can detect effects will vary across cohorts–we illustrate them here for a specific case. These findings provide a possible way to reconcile the disparate findings in the literature.

5 Estimating the baseline model

We now estimate our model and then study the long term effects of shocks occurring early in life. We start by providing some details on how we estimate the model and then we estimate the model for two populations that are in close-to-stationary environments. We then study the effects of temporary shocks and estimate the effects of WWII on women's lifetime mortality.

5.1 Simulated methods of moments

To estimate the parameters for a given cohort we use the simulated method of moments.[20] We choose a starting value for the parameters,[21] simulate mortality rates and compare the simulated rates to the ob-

[20] Because we observe only the mortality – and not the distribution of health – for each cohort every year, we cannot use simulated maximum likelihood methods. There is no empirical counterpart to the distribution of H_t, only the fraction that is below the threshold.

[21] We choose the initial values as follows. An initial guess for the level of initial health is giving by infant mortality. If infant mortality is 15%, then in the absence of shocks and investment, the mean initial health must be one standard deviation away from the threshold. So a value of $\mu_0 = 1$ provides a reasonable starting point. Individuals typically double in size within the first year, so we can guess that I could be half of μ_0. We set the standard deviation σ_ϵ to be equal to I. The depreciation rate for various stocks is usually estimated

served rates. We then iterate until we find the parameters that can best predict the data, that is those that result in the smallest prediction error. More precisely we chose the parameters by minimizing the sum of sqaured errors between actual survival and predicted survival at each age:

$$\min_{\theta} \left[SR_t - SR_t\left(\theta\right)\right]' W \left[SR_t - SR_t\left(\theta\right)\right]$$

where $\theta = \{\alpha, \delta, I, \mu_H, \sigma_\epsilon\}$ and W is a weighting matrix–we set it to be the identity matrix in our baseline estimates and check if the results are robust to alternative weighting schemes. We impose that σ_ϵ, δ and α be non-negative, but otherwise make no parameter restrictions. To guard against the possibility that the initial guesses will determine the outcome or that the algorithm stops at a local minimum we implement the procedure suggested by Powell (1964) in addition to using MATLAB's fminsearch. Details are in the Appendix.

We target the survival curve. This is a commonly used criteria in epidemiology and the results are easily interpretable: we can summarize the fit of the model based on how far the predictions are from observed life expectancy. Ideally the objective function that we choose to target for estimation would not affect the estimated parameters. But this is only true if the model is correctly specified. So we also report the fit in terms of the log of squared errors of the log of mortality, which we have checked graphically (in the appendix we also show how the model fits alternative criteria).[22]

Although standard errors can be bootstrapped, we do not report them here. Because each curve is traced out from large populations the standard errors are effectively negligible, as as been noted elsewhere (Honoré and Lleras-Muney, 2006). The most important source of error will be model specification error as the estimations will show.[23]

Lastly there are some important measurement issues. Migration flows are likely to add noise in our estimates. The HMD data is taken from vital statistics, which count the number of deaths occurring in France, including migrants who died on French soil and excluding deaths of French individuals migrating out of France. The population counts are not corrected for in and out migration either. As a consequence, the true data generating process is likely to be a mixing distribution of cohorts of the same age but born in different countries. We cannot address this issue well–there is no cohort data that would allow us to seriously correct our estimates for migration. We leave this to future work. Also population and deaths are poorly counted during wars, and particularly when there are changes in territory. There are in fact many changes in territory (see Appendix), but our fundamental conclusions did not change much under various assumptions to deal with this. Finally we use data up to age 100, but the HMD has estimated the numbers after age 90. We also check for the robustness of the results when we limit our data to age 90.

to be a few percentage points. but in our case the depreciation grows with t, so we assume that δ is 0.001 to start with. And we set α=1, assuming the depreciation grows just linearly with time. The results presented here however use starting guesses that have been updated many times.

[22]We could look at other objective functions. Life expectancy weights early ages very heavily because early deaths result in large losses in terms of years lived. Conversely it weights mortality in old ages less. If we minimize the error in the level of mortality, we give equal weight to errors at all ages and penalize large deviations in levels (at very young and very old ages). If we minimize the errors in logs, then we minimize deviations in percentage terms–this effectively gives more weight to errors that are large relative to baseline, so in effect it weights reproductive ages more heavily. Ultimately the choice matters if the model is mis-specified.

[23]We have nevertheless bootstrapped the standard errors and verified that they are very small compared to the changes in the parameters that are observed when we modify the model. Results available upon request.

Figure 9: Survival curve for apes

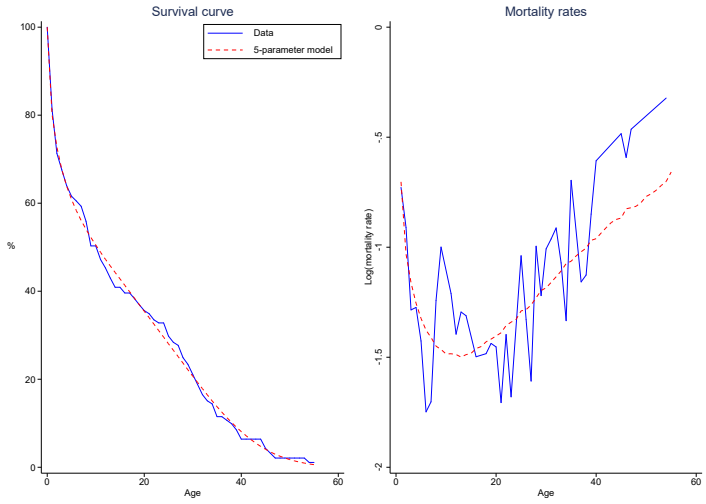

Note: The figure shows the estimates for the basic 5-parameter model without lifetime accidents. It also does not model the adolescent hump. The estimated parameters are in 1.

5.2 Estimates for primates.

The model is not unique to humans and so far simply describes biological disease and aging processes in a stationary environment. Several prior studies suggest that primates and human share common biological aging patterns and they have similar mortality profiles, particularly in old age (Kohler et al. 2006, Bronikowski et al. 2011). In addition to their biological similarity to humans, primates are also of interest because they likely face a stationary environment and they cannot reallocate resources over the lifetime: they have no access to saving, borrowing or medical technologies. Thus they provide a good test case for the basic model which assumes constant parameters, including resources, at every age.

We concentrate on chimps, who are the closest primates to humans. Data on chimps living in the wild is taken from Bronikowski et al. (2011). These are collected from surveillance sites that track the composition of groups over time. These data have some advantages: the populations are relatively large (we have about 80 observations) compared to what is available from animals in captivity, they include mortality at all ages and are available by gender and species. However because animals migrate in and out of a given group and geographic location, there is larger measurement error.

The results from the estimation are in Table 1 and Figure 5.2. Both show that our model fits these survival rates remarkably well. The actual life expectancy for females is 15.38, while our predicted value is 15.35. Thus our model provides a good description of the evolution of mortality for chimpanzees. The model also fits mortality rates well, but the mortality rates are very noisy, so there remains substantial uncertainty about the parameters. So we now consider whether the model fits large human populations.

5.3 Baseline cohort in a close-to-stationary environment

We now estimate the parameters of the model for the 1816 cohort. Period and cohort life expectancy are very similar for this cohort suggesting an environment that is also close to stationary for this population (life expectancy for this cohort was 41 based on period tables, and about 40 using cohort tables, see Figure 8).

Figure 10: Lifetime mortality for French women born 1816

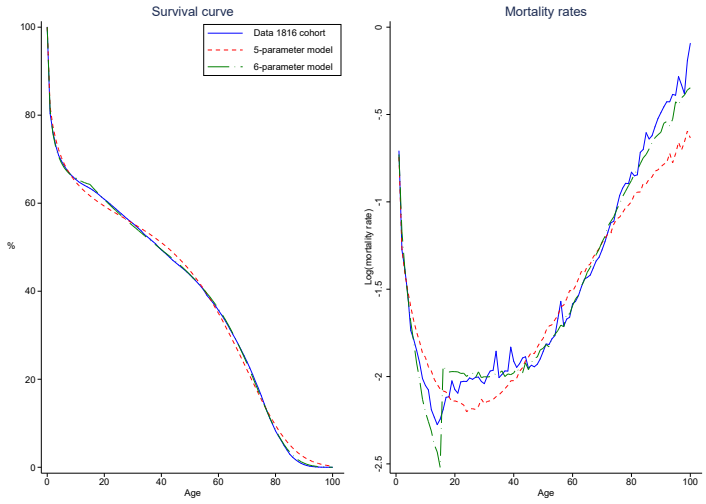

Note: The 5-parameter model assumes no accidents occur throughout the lifetime. The 6-parameter model allows for a positive accident rate to start in adolescence. See 2 for parameter values.

Figure 10 shows the results of estimating the baseline 5 parameter model, and table 2 shows the estimated parameters and the fit (see column 1). We provide three measures of the fit: how well we match the survival curve (the object we target), the log mortality rates, and the distribution of the age at death.

The basic 5-parameter model predicts life expectancy at birth of 38.43 which is very close to the actual life expectancy we compute of 38.25. But the figure shows that the fit, particularly in terms of mortality rates, is among the oldest and particularly around adolescence and young adults. So next we use the model to investigate how to model the adolescent hump.

5.4 Understanding the adolescent hump: the effects of permanent changes occurring after birth

To understand how to best model the adolescent hump, we simulate the impact of a permanent shock occurring at age 12. The changes in hormonal levels occuring in adolescence have large biological and behavioral effects. Puberty also marks the begining of adulthood and in many societies marriage, work and living arrangements change substantially as a result.[24] We consider five different models of what occurs in adolescence: a decrease in the annual investment level, an increase in the accident rate, an increase in the variance, an increase in the depreciation rate and, finally, an increase in the death threshold.

Figure 11 shows that these changes all increase mortality. But they affect the age-profile of mortality in different ways. Investment declines and depreciation increases both permanently raise mortality for all subsequent ages, although the effect of depreciation is barely noticeable at first it cannot explain the adolescent hump.[25] Accident increases generate a floor between ages 12 and 40, and then the profile converges

[24] Adolescence is a period of great change in many dimensions. For a review see Dahl, Ronald E. et al. (2018).

[25] Case and Deaton (2017) observe that for cohort born in the US after 1950 mortality rates in adulthood are becoming progressively higher. These simulations suggest that cohorts born after 1950 experienced deteriorating health and mortality profiles because they have lower health resources or higher depreciation. Because more recently born cohorts are born in better health (and experience lower infant and child mortality up to age 20 (Currie XXX) the the model further suggests that a worsening of conditions starting at age 20, such as deteriorating labor market opportunities, is responsible for these deterioration. See Lleras-Muney (2017) comment for further details.

Figure 11: Effect of permanent changes at age 12 on (the log of) mortality

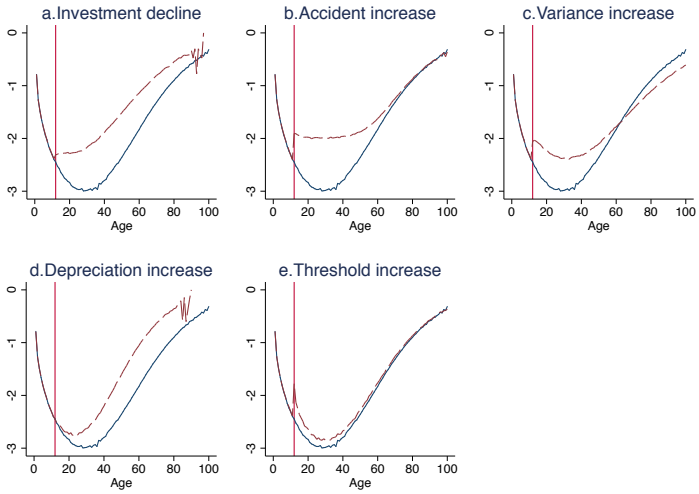

Note: All increases are 50% increases, except for the threshold which is increased to the level of μ. The baseline parameters are the same as in Figure 3

to counterfactual. Variance increases also generate something close to a floor but the affected cohort's mortality exhibits a cross-over: it has lower mortality than the original cohort after some age. Finally threshold increases result in a peak and then the profile slowly converges to counterfactual. This peak is similar to what is observed in wars, which we consider below, but not that similar to the adolescent hump.

These simulations suggest that this bump can be best conceptualized in our model as either a variance or an accident increase starting in adolescence. But decreases in I or in the threshold could potentially also rationalize this hump. Therefore we test four models for the adolescent hump in the data next.

5.5 Estimating the baseline model accounting for the adolescent hump.

We fit the baseline model again allowing for a permanent shock to occur starting in adolescence. Adolescence (puberty) is assumed to occur when girls have their first menstruation.[26]de La Rochebrochard (2000) reports that the onset of menarche occured around 15.8 in 1816, so we assume it starts at 16.[27] We estimate 4 models: increasing the accident rate, increasing the variance, increasing the threshold or decreasing investment.

The results in Appendix Table 2 show that while all types of shocks improve the fit relative to the baseline model, accidents improve the fit the most (the fit of the survival curve goes from 155 to 12). The parameter estimates illustrate that not modeling reproductive-age mortality not only lowers the fit, it biases the parameter estimates, and changes life expectancy predictions. Figure 10 illustrates the results for the best and final model. It shows that by modeling the adolescent hump we also correctly predict mortality among the eldest. We predict a life expectancy of 38.28 whereas the actual life expectancy is 38.25. Visual inspection of the fit in terms of log mortality rates (or the distribution of the age at death) also suggests an excellent fit.

[26]The onset of puberty could be conceptualized as occuring earlier, when the the adolescence growth spurt starts, heights increase, and breast and other secondary sexual traits start to appear. But historical data for these alternative measures are not available.

[27]The age at menarche (first menstruation) in France fell from around 16 in the second half of the 18th century to about 13 in the second half of the 20th century. The best estimate for the age at menarche over time is given by: Age at menarche = (- 0.0175 x calendar year) + 47.4. Using these estimates we predict that the onset of menarche was 14.85 in 1860. See de La Rochebrochard (2000).

But these results are based on a single cohort. To verify these findings we re-estimate the model for the 1860 cohort. For this cohort the assumption of stationarity is less sensible: cohort life expectancy substantially exceeds period life expectancy at birth suggesting large changes in the environment. Thus the assumption that the parameters are constant throughout the lifetime is more likely to be violated. However Table 3 shows that the conclusions about how to best model the adolescent hump for females are the same: the model that allows for a positive accident rate that increases at puberty provides the best fit.

Using these models we compute what lifetime mortality would have been in the absence of the adolescent hump–which has mostly disappeared in recent cohorts.[28] Excess mortality starting in adolescence is estimated to have lowered female life expectancy by around 7.5 years for the 1816 cohort and by 7.8 years for the 1860 cohort.

It may seem surprising that reproductive age mortality is best captured as an increase in the accident rate, and therefore not related to health status. However this finding is consistent with the findings in the literature. Two causes of death have been documented to account for a large fraction of deaths among women ages 15-45 in non-war times: maternal mortality, and "external" causes, which include traffic accidents, poisoning and violent deaths (including suicide and murder). Loudon (2000) argues that historically poor hygiene and obstetric practices were mostly responsible for infections (sepsis) and hemorrhage–the main reasons why women died during childbirth. These poor obstetric practices were widespread, so maternal mortality was large and it killed both rich/healthy and poor/unhealthy women.[29] Other external causes of death appear to also be unrelated to health status. Finally we note that while there is no data on causes of death by age and gender for 19th century France, contemporary data shows that mortality rate from external causes of death are well approximated by a step function. This is shown for the case of the US in Figure 12.

We focus on our preferred 6-parameter baseline model (in column 4 of tables 2 and 3) to discuss the estimated parameters. In 1816 the initial health is about 0.86 standard deviations away from the threshold. This suggests that absent any shocks or any investment in the first period, infant mortality would have been roughly 15% (instead it is 17%). The annual investment is about 0.4 of a standard deviation, so in the absence of shocks, the health stock is increasing by roughly half in the first year of life, and doubling by the second year. The variance of resources is about 1 standard deviation, roughly equal to the variance of the initial stock of health. The annual depreciation rate δ is very small and equal to 0.0006 of a standard deviation. But the rate is increasing over time exponentially: α is around 1.8. Finally we estimate that the adolescent accident rate is about 9 per thousand.

Comparing 1860 to 1816 we find that the initial stock increased from 0.86 to 0.93. The mean level of resources fell, but so did the variance of resources. Both aging parameters fell, but most significantly the depreciation rate fell from 0.0006 to 0.0004. Finally the accident rate also fell substantially from 9 to 7 per thousand. These comparisons suggest that the environment is improving substantially, consistent with the observation that life expectancy rose by about 5 years between these two cohorts. But the tables also suggest caution. The parameter estimates are sensitive to the specific model we estimate. For example adolescence could also reasonably be modeled as an increase in variance (this model has the best fit in terms of the distribution of the age at death). Comparing columns 3 and 4 in either table we see that the estimates for all parameters differ substantially across models, with the exception of the estimates for μ. For instance for

[28] In this counterfactual we set the accident rate to the the same all throughout life.

[29] The disappearance of the hump described in Figure 1 is also consistent with the elimination of maternal mortality as a cause of death after the 1930s. In France maternal mortality between 1850 and 1890 is estimated to be around 5 per 1,000 births (Bardet et al. (1981)) and to have remained at that level until the 1920s, it fell rapidly after the mid 1930s (Loudon 1992). By 1970 maternal mortality was around 28 per 100,000 (Bouvier-Colle et al. 2008)

Figure 12: US Age-specific Mortality rates per 1,000 in 1990, by age and cause of death

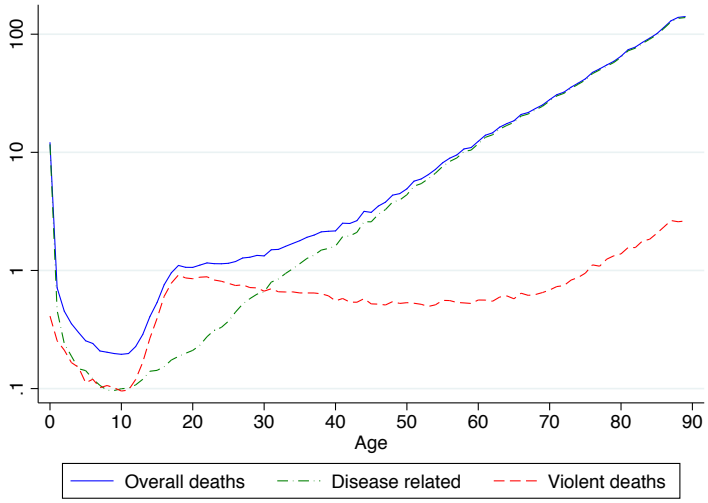

Note: this figure is reproduced from Schwandt and von Wachter (2018)'s paper "Mortality Profiles of Unlucky Cohorts: Effects of Entering the Labor Market in a Recession on Longevity" who generously agreed to let us use it. The data come from period (not cohort tables) so they are not directly comparable to ours but we use it to demonstrate that the mortality rate from non-disease related causes of death is well approximated by a step function turns on in adolescence.

the 1860 cohort I is estimated to be 0.4879 in column 3 but it is 0.3318 in column 4, so it is 30 percent lower, which is a very significant change.

5.6 Robustness and sensitivity checks.

In Table 4 we perform additional checks on this basic 6-parameter model using the 1816 cohort. The fit is still poor around the time of the onset of maturity. This is in part the results of two assumptions: that adolescence's onset occurs exactly at 16 for the entire population, and that there are no accidents before age 16. In column 2, we show the estimates if we assume that the accident rate is positive before adolescence and increases at age 16. The fit of this 7-parameter model is not better than the fit of the 6-parameter model, despite the added parameter.

Next we allow for the onset of maturity to be normally distributed. Unfortunately there is no data on the distribution of the age of menarche for 1816, only its mean. So we make use of the 1970 distribution (again provided in de La Rochebrochard, 2000) to obtain the standard deviation. This augmented model results in further improvements in the fit (relative to column 1) but we do not use it as our baseline model because there is no good data on the standard deviation for the early cohorts we study. Alternatively we can assume that the onset of menarche occurs at a normally distributed time and estimate the parameters of this distribution. Not surprisingly this improves the fit but notice that the estimated mean and variance are far from the observed ones, so we do not use this model in the remainder of the paper.

In column 5 we investigate what happens if we use the (normalized) number of deaths as weights in the estimation. In column 6 we use weights and target the distribution of the ages at death instead of the survival curve. This considerably worsens the survival fit and does not improve the fit of the distribution of the age at death by much. In the last column we use only data up to age 90 to see what the effect of censoring is and because the data after 90 are estimated. The estimates are somewhat sensitive to these

Figure 13: Effect of exogenous temporary shocks at age 20 on log mortality

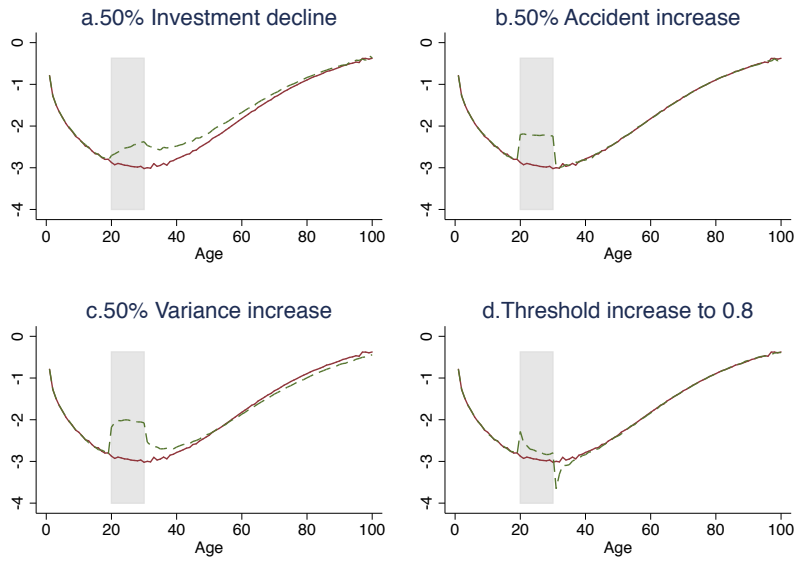

Note: Simulations for two populations of 500,000 individuals each. The shaded area corresponds to the years of the temporary shock. The baseline parameters are the same as in Figure 3.

choices. The predicted life expectancy is very close in all cases (the error in the predicted life expectancy is less then 0.1 years of life), except when we target the age at death distribution (the prediction is off by about a year). But the counterfactual predictions are sensitive: eliminating the hump results in a loss of life of 7.58 years on the baseline model and 9.81 in the worse model. These results suggest several conclusions. First the model could be improved with additional data on the distribution of the onset of maturity and with data on the distribution of health at various ages. Second model specification and estimation procedures matter for the estimates. Lastly some features of the data remained unmodeled. For example for 1860 there is a visible spike before the adolescent hump (corresponding the the 1870 war) that we have not modeled. This unmodeled shock will generate bias in the estimates. We investigate these next.

6 The effects of temporary shocks on health and mortality

We now use our baseline 6-parameter model to understand how temporary shocks, like wars, recessions or infectious disease epidemics, affect the profile of health and mortality. We start by simulating the effect of temporary changes in parameters starting at age 20 and lasting 10 years (for this simulation we are ignoring the adolescent hump for clarity).

Figure 13 illustrates the effects of each type of shock. Each type of shock leaves a unique imprint on the mortality profile of the affected cohort. Temporary decreases in investment levels generate spikes in mortality, similar to those observed for wars. When investment falls, mortality rates start to rise, they peak the last year of the shock and they fall back thereafter. But mortality rates remain elevated throughout the lifetime (relative to the counterfactual of no shock) thus generating "scaring." By contrast, temporary increases in the accident rate immediately increase mortality but have no permanent effects: mortality goes back to its initial path immediately after the shock ends.

Increases in the variance result in a sustained increase in mortality during the shock period. But after

the shock ends, mortality falls below counterfactual levels — this is the result of having individuals with large positive shocks. But both variance and accident increases have similar effects.

Finally, an increase in the threshold appears to generate "harvesting"— earlier deaths for the frail. It results in very high mortality in the first year of the shock. But mortality starts dropping before the shock ends. Once the shock ends, it dips below counterfactual mortality and then rises back up and converges to its counterfactual level. This is because all frail individuals are killed when the threshold first increases. And when the threshold is restored to its original (lower) level mortality falls substantially because there are very few individuals close to the threshold. This pattern fit the effects of extreme weather or pollution events, which appear to displace deaths in short term.[30]

6.1 The effects of WWII

Wars correspond to the largest temporary change in mortality rates over time for individuals of reproductive ages, leaving visible spikes in the mortality profile. WWII is the longest conflict in our sample lasting six years–this should make it easier to distinguish among different type of shocks.[31]

The simulations above suggest that only investment shocks can generate patterns that match those we observe in the data, with mortality rising every year of the war, and peaking in the last year. But other types of shocks also generate spikes. So we estimate the structural parameters explicitly allowing for a shock lasting six years, and varying the type of shock. We first concentrate on the 1921 cohort, who turned 18 when WWII started in 1939. We start with this cohort because it has complete mortality up to age 90, whereas younger cohorts are censored. A disadvantage of looking at this cohort is that the assumption of stationarity is likely to be violated, since cohort life expectancy was increasing substantially throughout the 20th century.

Appendix Table 5 shows the results. We evaluate the models in terms of how they fit the overall profile of survival and whether they fit the shock itself. Surprisingly, given our simulations, all models provide almost equally good fits of the survival rate (or the log of q) and the predicted life expectancy. However the fit during the war is clearly better matched by the model that assumes WWII was equivalent to a decline in I. Although all models underestimate the mortality rate and the number of deaths during the war, the I-shock makes the smallest mistakes. It underestimates the number of deaths by 21%. By comparison a pure accident model underestimates the number of deaths by 36%, a variance model by 45 percent and a change in the threshold by 32%.

Figure 14 shows the fit for this model. We estimate that the war lowered life expectancy by approximately 5 years for the 1921 cohort. This of course includes a large number of deaths that occured during the war. But we can also assess the amount of scaring. Conditional on surviving to 1945, life expectancy is 5 years lower than it would have been in the absence of the war–this is very large. As the table shows the predictions for other shocks are starkly different: an increase in the threshold during the war generates pure selection, and since the weakest have been killed off, mortality rates are predicted to be lower after the war and life expectancy higher conditional on survival.

We conclude the war is best characterized as a decline in health resources I for women.The parameter estimates show a I moving from a lifetime value of 0.29 to a value of -0.11 during the war years.[32] This is

[30] For both pollution and weather there is evidence of both short term displacement and longer term effects on mortality. For instance see Schwartz (2000) or Zeger et al (1999) for the effects of pollution. For the effects of very hot or very cold weather see recent articles by Deschenes and Moretti (2009) and Deschesnes and Greenstone (2011).

[31] It was also very intense. WWII is estimated to have killed around 600,000 individuals in France, about 1.4% of the 1939 population.

[32] We conduct an additional exercise to verify these findings: we compare the the death profile of French women with that of Swedish

Figure 14: WW2 as an investment shock. French Women born in 1921.

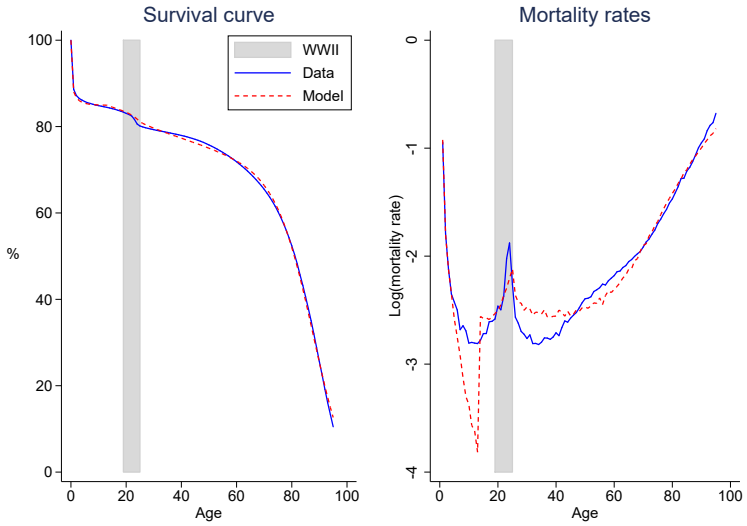

Note: see 5 for parameters value. This model assumes a decrease in the investment level every year of the WWII.

consistent with evidence showing that GDP, food supplies and sanitary conditions declined substantially during the war.[33] Infant mortality rates, which are very sensitive to these inputs, rose substantially during this period.[34] Interestingly recent work also finds results consistent with our findings that WWII generates scaring. Kesternich et al. (2014) study the effects of the WWII across 13 European countries. They find that individuals more exposed to the war experienced worse economic and health outcomes later in life than other survivors who were less exposed. Havari and Peracchi (2017) and Schiman et al (2017) report similar findings for WWII. This result however does not imply that all wars are necessarily equivalent to declines in I.[35]

There are a few limitations to these results. First the population and death counts during the war are particularly poorly measured. There were large changes in territory during the war years and substantial migration. So the population and death numbers are subject to substantial measurement error.[36] Second the 1921 cohort is only followed until age 94–if the war changes mortality rates thereafter, we are not properly computing the effects on life expectancy because of censoring. Third, we assume that the shock started in 1939 and ended in 1945, and that it can be modeled as an equal decline in I every year the war. However

women. Sweden was the only country in the sample that did not participate in the war. Appendix Figure 24 shows that compared to Sweden, mortality rates in France are substantially more elevated during the war years, and they remain elevated for many years after the war, similar to the simulated effect of an investment decline. However these comparisons are imperfect because Swedish mortality rates before the war (up to age 19) are different than the French's. According to our model this would result in different mortality rates after age 20 even in the absence of the war, so the swedish case provides only suggestive evidence on this question.

[33]GDP declined substantially during the war. Moreover Occhino et al. (2006) estimate that between 20 to 55% of GDP was appropriated by Germans every year of the occupation. There was a substantial decrease in the availability of food–food rationing began in 1940. There was also a deterioration in sanitary conditions in France. For example diphteria cases among school-aged children rose per 100,000 increased from 32.3 (in 1940) to 110 in 1943 (Stuart 1945).

[34]In the HMD mortality in the first year of life was 0.063 in 1938 an declining. It rose to a high of 0.085 in 1940, the worse year of the war.

[35]But other wars also appear to have caused scaring. For instance Costa (2012) documents that surviving soldiers in WWII have higher morbidity and mortality later in life. This evidence pertains only to male soldiers however. Other studies have documented scaring effects of war in utero–for a comprehensive review of the literature on in utero shocks see Almond and Currie (2011). More recently Lee (2017) presents evidence of substantial health effects in adulthood of the Korean War for cohorts exposed in utero. Estimates in this section pertain to adult females.

[36]See appendix notes for how we treat the data during these years.

Figure 15: The effect of the 1918 flu pandemic on mortality for the 1900 cohort

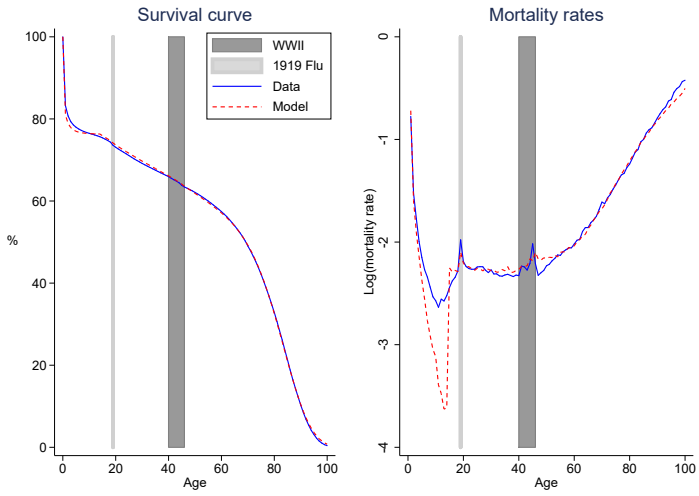

Note: Both the 1918 flu and WWII are modelled as a temporary decrease in I.

rationing continued until 1949 so the decline in resources might have lasted for longer. And not all years were equally difficult. The invasion which took place in 1940 was by most accounts much worse than the year before or after (based for instance on what occurred to overall mortality rates). Fourth, we choose the most parsimonious model to fit the data using a single parameter–it is possible for instance that the war also affects the depreciation rate. Finally these results assume there were no compensatory responses to WWII. But there were efforts like the Marshall plan to help rebuild infrastructure and promote economic growth which started in 1948. We investigate these in the last section of the paper.

6.2 The Effects of the the 1918 Flu pandemic

We next estimate the effect of the 1918 flu pandemic, by far the largest infectious disease epidemic in the last 200 years in France occurring at the end of WWI. We focus on the 1900 cohort which turned 18 in 1918. We model WWII as a change in I based on our previous results and test which temporary change in parameters fits the 1918 shock best. Because the cohort data do not clearly show increases in the 1914-1917 period, we simply model the effect of a one year shock (see Figure 15).

The results are in Table 6. Again we find that changes in I provide the best fit for survival, and it has the best predictions for the number of deaths and the mortality rate in 1918. The results for this model are displayed in Figure 15 for the flu pandemic. The combination of the 1918 flu and WWII results in a loss of life of 2.12 years, roughly 2 from the flu and only 0.2 as a result of WWII, which this cohort experienced at age 39.[37]Like in the case for WWII, we find that conditional on survival to 1919, scaring causes a large amount of scaring among the survivors, who live 2 years less as a result of exposure to the flu.

The results also suggest caution–once we account for the flu, the fit for WWII is poor and the model under-estimates the deaths from WWII. All models over-estimate deaths and death rates during the flu years. This illustrates the difficulty in estimating these effects without having information on the shocks themselves or additional data on health. This is particularly true in the case of the flu because, unlike the

[37]The effects are not additive–there are interactions.

Figure 16: Effects of a temporary decrease in investments in childhood on health and mortality, by age at the time of the shock

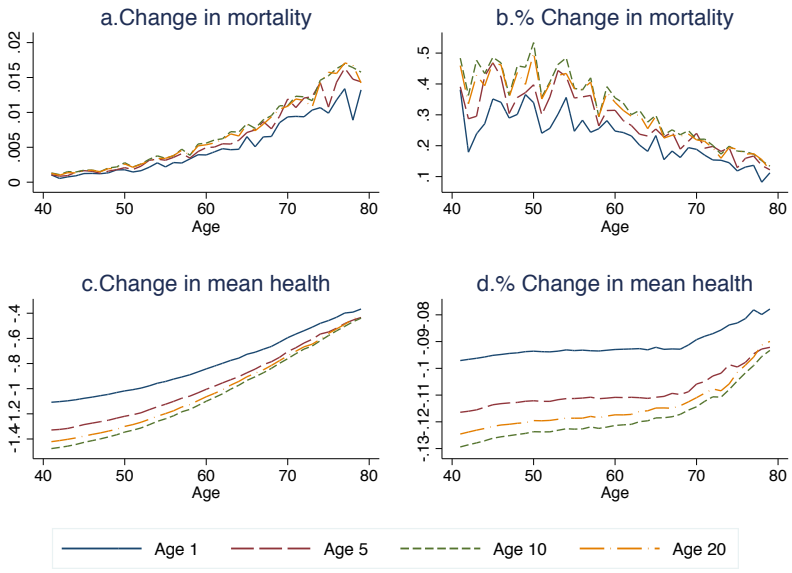

Note: Simulations for two populations of 500,000 individuals each. The shaded area corresponds to the years of the temporary shock, which is a 50% decline in I for 10 years. The baseline parameters are the same as in Figure 3 but we add an adolescent hump at age 14 with $\kappa = 0.001$.

war case, it is a one period shock which makes it harder to identify which model is the best–as table 6 shows, all models provide reasonable fits except for the variance shock.

Nevertheless the results are consistent with findings from previous literature (Almond, 2006, and Beach et al. 2018), although this literature has concentrated on in-utero effects, whereas we are documenting effects of insults during prime age years. In the next section we investigate how I-shocks affect cohorts based on the age of onset of the shock.

6.3 Effects of a temporary I-shock at different ages in childhood on health and mortality in adulthood

We found that WWII and the 1918 flu pandemic are best characterized as temporary decreases in I. We now investigate the effects of an identical I-shock occurring at different ages on adult mortality and health to match what is typically done in the literature. We simulate a 50% decline in I for 10 years starting at ages 1, 5, 10 and 20, using our 6-parameter baseline model, with a hump at age 14.

The results of the simulations are in Figure 16. It shows the change in mortality and health that is experienced by different cohorts, relative to a baseline without the 10 year shock.

The results show several interesting patterns. First mortality rates for affected cohorts are elevated for all ages in adulthood, but the effects on mortality are very small at first. In levels these effects increase with age in adulthood, but in percentage terms the effects fall with age (because mortality rates are rising after age 40). Second health effects decline with age (in levels and percentages). Third the largest effects on health and mortality are experienced by those who were 10 at the onset of the 10-year shock, not among those who were 1 or 5. This is occuring because the mortality of the 1, 5 and 20 year olds is larger than the mortality rate of 10 year olds, so some of these individuals would have died even in the absence of the shock. Of

course these simulations only indicate patterns for a given set of parameters–we have provided no proof that these patterns are always true. But the results do sugest that those affected earlier need not have the largest longtime impacts, when these are measured at a point in time later in life.

These simulations assume that the shocks are equal at all ages. But this is not clear: for example WWII might have affected 20-year olds more then 10-year olds if the women were involved in the fighting. There are other limitations to these results. For example they assume that the shock did not affect the onset of adolescence. But the literature on height for instance suggests otherwise: poorer populations grow slowlier and for longer. Despite these limitations, these simulations demonstrate the usefulness of the model in providing a clear account of how early shocks affect the entire profile of health and mortality.

7 Implications for optimal investments

7.1 Optimization in a stationary environment

For now we have assumed a constant investment profile over the lifetime. But would that be an optimal allocation of resources over the lifetime? In this section we show that a social planner concerned with maximizing the life-expectancy of a population would choose an investment profile that ultimately results in patterns of mortality with striking similar shapes of the ones studied in the previous sections. In other words, the optimal investment sequence does not fundamentally change the shape of mortality.

First we develop notation to describe the problem that would face benevolent social planner. We solve this problem under 2 key assumptions. The first key assumption is that the planner has a fixed budget but has the ability to borrow and save costlessly–on other words the planner knows exactly what the total lifetime resources are for a give cohort and these resources can be redistributed across the lifetime at no cost. The second assumption we make is that the planner wishes to maximize life expectancy.

The survival function tracks the probability of surviving over time. It is naturally expressed as a function of the cdf of health in the population. The probability of surviving until the end of period t is $S_t = 1 - F_t(0)$. Life expectancy at birth is conveniently related to the survival function

$$LE = \sum_{t=1}^{\infty} S_t$$

Several observations are in order. First, in practice, this is a finite sum. Second, contrast this concept with the "period" life expectancy usually computed. If the distribution where stationary over time, then the two concepts would coincide. But as the data shows and our estimates corroborate, the mortality rates are not stationary.

Now suppose that instead of keeping I constant that the social planner can choose an investment path $\mathcal{I} = \{I_t\}$ that is age-dependent. Also assume that the budget (B) over the lifetime is fixed but that the planner can move resources over time periods costlessly, as if a perfect annuity were available.[38] Then the optimization problem takes the form

[38] This is a standard assumption in this type of models, for example see Murphy and Topel (2006).

$$\begin{aligned} \max_{\mathcal{I}} LE(\mathcal{I}) &= \max_{\{I_t\}} \sum_{t=1}^{\infty} S_t(\mathcal{I}) \\ s.t. & \sum_{t=1}^{\infty} I_t \cdot S_t(\mathcal{I}) \leq B \end{aligned}$$

The first order conditions for this maximization are given by:

$$\sum_{s=t}^{\infty} \frac{\partial S_s(\mathcal{I})}{\partial I_t} - \lambda \left[S_t(\mathcal{I}) + I_t \frac{\partial S_t(\mathcal{I})}{\partial I_t} + \sum_{s>t}^{\infty} I_s \frac{\partial S_s(\mathcal{I})}{\partial I_t} \right] = 0$$

where is the Lagrange multiplier and it represents the shadow cost for the social planner of an additional year of life expectancy. All the terms in the bracket are positive.

The FOCs imply that on the optimal investment path, the marginal effect of increasing investment at a given age must be equalized across all ages. Increases in life expectancy (the first term on the left-hand), must be balanced by the losses incurred by having to tighten the budget at subsequent periods to keep the budget balanced (the term in brackets).

7.2 Timing of optimal investments, polynomials

A full nonparametric approach for the optimal investment profile over the lifetime would require optimizing over a hundred or so parameters (one for each age) for each cohort. In the absence of a closed-form solution, this is impractical. It is also not feasible since we have 100 data points: if we allow for a unique investment level at every age we are under-identified (we would have 100 data points and 106 parameters to estimate). Instead, we follow a lower-dimensional sieves estimation method.

We start with by approximating the investment profile over age with a second order polynomial which adds only 3 more parameters to the estimation. We impose the constraint that the total spending per cohort is the same as in the constant investment case. Given a budget B we run a grid search to find the quadratic investment profile that maximizes the life expectancy of the cohort.

The results of this exercise are in Figure 17. We find that a U-shape investment profile is optimal to maximize the average life-expectancy in the population (panel a). Notice that although our original model sets I to be constant in levels, in percentage terms, relative to the baseline level of health at a given age, I was already U-shaped. What we find is that the optimal investment is even more U-shaped – that is, it transfers additional resources to the young and the old, away from the middle-aged individuals. Interestingly health care expenditures by age in most countries actually follow this age-profile (Alemayehu and Warner, 2004). In the specific case we show in Figure 17, based roughly on parameters for French women born in 1870, optimizing investment results in a gain of about 3 years of life expectancy. Panel b shows the mortality curve before and after optimization–it has the same basic shape we have observed. These results show that optimal health investments are largest when health is at its lowest, that is at very young and very old ages. This is consistent with empirical findings which show that health and the demand for medical services are negatively correlated (Wagstaff, 1986).

7.3 Other Properties of investments

Each investment profile, $\mathcal{I} = \{I_t\}$, generates a sequence of distributions of health, $F_{H_t}(\mathcal{I})$, and its associated mortality rates, $MR_t(\mathcal{I})$. How are investment decisions at different ages related? We show in the following

Figure 17: Optimal Investment Levels by Age

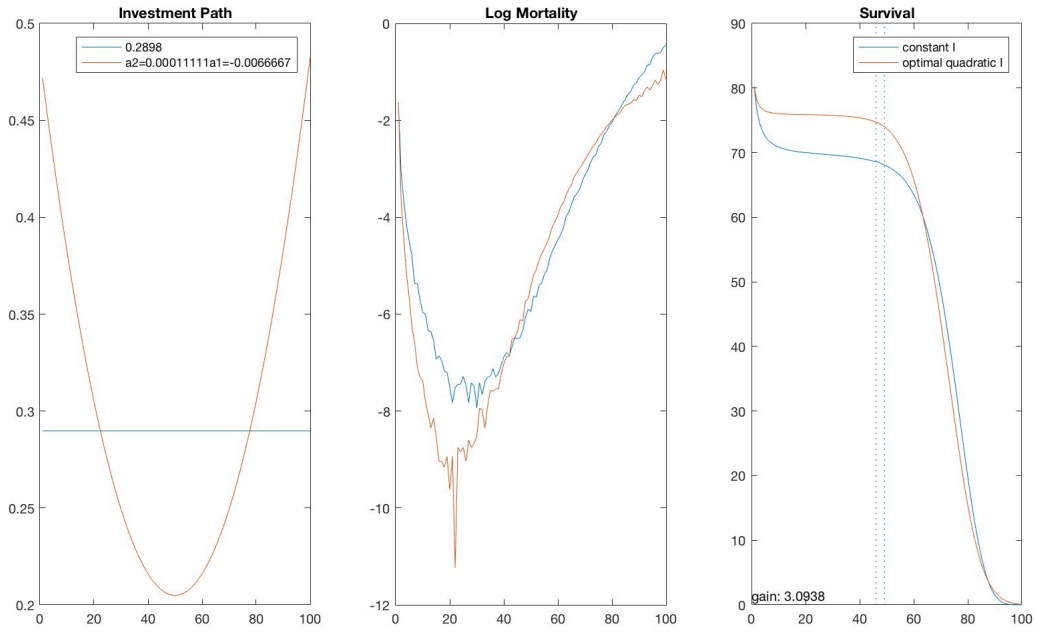

proposition that investments in health are dynamic complements.

Proposition 4: Investments exhibit dynamic complementarities Optimal investments in health are dynamic complements throughout lifetime. $\forall t_1, t_2\ t_1 < t_2$, $\frac{\partial^2 MR_{t_2}(\mathcal{I})}{\partial I_{t_1} \partial I_{t_2}} \leq 0$. See proof in the appendix.

The complementarity arises through the health accumulation process. A higher investment at time t pushes the distribution of health up moving more individuals away from the threshold. Additional investments make it more unlikely that negative shocks will push individuals below the threshold. Although the mechanism is different this result is similar to the results in Cunha and Heckman (2007),who study optimal investments in human capital in a model without attrition (mortality).

This result has implications for optimal compensation: if a cohort suffers from an unlucky shock and the planner wishes to compensate them so that survivors can enjoy the same mortality rates they would have in the absence of the shock, how can the planner achieve this, and how does this vary with the level and timing of the shock? Our results imply that negative shocks need to be more than compensated for, because of complementarity. In other words a decrease in investments needs to be followed by an increase resources that is greater than the loss, to give the survivors the mortality profile they would have experienced in the absence of a shock.

7.4 Optimization when budgets depend on health.

We have solved the optimization problem under the (strong) assumption that resources are not a function of population health. But if food and other resources are produced rather than taken from the environment, health is likely to impact resources by affecting the work capacity of the population. Indeed nutrition levels and disease rates have been shown to affect productivity and wages (Thomas et al., 2004). They also affect

inputs into wages such as cognition and education (Field et al., 2009). Many empirical studies report a correlation between income and health (Cutler et al., 2012, Chetty et al., 2016). While our baseline model embeds the effect of resources on health, a causal link going in the other direction is also at play: people who get sick or are hospitalized suffer a subsequent drop in income (Smith, 1999, Dobkin et al. 2018).

The simplest way to incorporate this channel in our setting is to assume that people whose health lower than some disability or disease level, h_D , are unable to participate in production and thus generate zero income while people whose health is high enough generate income w .

$$\begin{aligned} \max_{\mathcal{I}} LE(\mathcal{I}) \quad &= \quad \max_{\{I_t\}} \sum_{t=1}^{\infty} S_t(\mathcal{I}) \\ s.t. \qquad & \sum_{t=1}^{\infty} I_t \cdot S_t(\mathcal{I}) \leq B = \sum_{t=1}^{\infty} w\left[1 - F_{H_t}(h_d;\mathcal{I})\right] \cdot S_t(\mathcal{I}) \end{aligned}$$

Again here we assume either an annuity market, or some pooling across cohorts in a stationary environment such that it is the cohort budget that matters, not the the within-cohort, per-period one. The first order condition is:

$$\sum_{s=t}^{\infty} s \cdot \frac{\partial S_s(\mathcal{I})}{\partial I_t} + \lambda \left[\sum_{s=t}^{\infty} \left\{ w\left[1 - F_{H_t}(h_d;\mathcal{I})\right] - I_t \right\} . \frac{\partial S_s(\mathcal{I})}{\partial I_t} - \sum_{s=t}^{\infty} \left\{ w \frac{\partial F_{H_s}(h_d;\mathcal{I})}{\partial I_t} + 1 \right\} \cdot S_s(\mathcal{I}) \right] = 0, \, \forall t > 1$$

or

$$\sum_{s=t}^{\infty} s \cdot \frac{\partial S_s(\mathcal{I})}{\partial I_t} = \lambda \left[\sum_{s=t}^{\infty} \left\{ w \frac{\partial F_{H_s}(h_d;\mathcal{I})}{\partial I_t} + 1 \right\} \cdot S_s(\mathcal{I}) + \sum_{s=t}^{\infty} \left\{ I_t - w\left[1 - F_{H_t}(h_d;\mathcal{I})\right] \right\} . \frac{\partial S_s(\mathcal{I})}{\partial I_t} - \right] \forall t > 1$$

First notice that $\frac{\partial F_{H_s}(h_d;\mathcal{I})}{\partial I_t} < 0$ for all t. (cf Proposition 1) and $\frac{\partial S_s(\mathcal{I})}{\partial I_t} > 0$. To estimate this model we need to make more assumptions than in previous settings. We need to estimate the threshold at which working capacity is positive (wages are positive). Ideally one would have data on wages over the lifetime for a given cohort and other institutional knowledge like whether there are laws regulating the employment of minors, in addition to data on mortality. For this reason we do not estimate this model. We leave this for future work.

8 Conclusion

This paper proposes a simple model of the evolution of health and mortality over the life course. This model is inspired by the basic observation that the age-profile of mortality is remarkably constant over time and cohorts. If health leads to mortality then we can learn about the underlying evolution of health by observing mortality rates. The basic model has six parameters and can be easily simulated and estimated. It can approximate quite well the mortality profile of cohorts born 1860 to 1940 and can be used to study the effect of temporary and permanent shocks occurring at different points in the lifetime. We demonstrate that in this model estimates of the effects of shocks depend on the age at which they occur. Moreover these effects have an age profile (they change mortality differently at different ages). This profile has different qualitative features depending on whether one measures them in levels, logs or some other transformation

of mortality rates, and they are different for health. Thus conclusions about long term effects of various shocks are sensitive to what metric (health or mortality) is used to measure them, when these consequences are measured, and the functional form used to estimate them. Thus disparate findings in the literature can possibly be reconciled by considering the differences across studies in these dimensions.

We use the model to estimate the effects of WWII and the 1918 flu pandemic on the mortality of survivors. We find that these events are best described as temporary changes in annual health resources. As a result we find scaring effects: those that survive are less healthy and their subsequent mortality is higher as a result of these shocks, despite the fact that on average these shocks killed the least healthy. Our findings are in line with previous empirical results but can be further used to estimate for instance optimal compensation profiles.

Our model places parametric restrictions on the evolution of health and mortality by age. A parametric model of health and mortality by age has many advantages. A long literature in demography and economics has struggled to separately identify age, period and cohort effects, which are not non-parametrically identified. Because age effects are parametric in our model, cohort and period effects can be separately identified. This is in fact illustrated here: we provide separate estimates by cohort, and we also estimate period effects like wars, separately for each cohort, and at different ages. However we do make some strong assumptions, for instance about the distribution of annual shocks, which should be further investigated.

This paper has some important limitations. First and foremost health is treated as an unobserved latent variable–we only demonstrate that our model of health delivers a mortality age-profile that is consistent with observed cohort mortality. Ideally one could use data on both health and mortality to better estimate the model. Secondly there is a scale and location normalization that cannot be avoided–there are 2 parameters of the model that cannot be identified. This makes it is difficult to interpret parameters across cohorts and countries: increases in initial health really could come from lower threshold or changes in standard deviation. Lastly our model does not provide closed form solutions and must be estimated using numerical methods. These methods are very sensitive to initial conditions and model assumptions.

But our model can be used to investigate many interesting questions that we have not considered here. We provide estimates across cohorts for women, but have not investigated gender differences or their evolution. We illustrate our methods for France, but these methods are also readily applicable to all other countries with high quality life tables, and could be used to investigate cross-country differences. The model can be extended to study correlations in health across generations in an over-lapping generations setting where the health (initial condition) of a new generation is made a function of the average health of mothers during reproductive ages. It can also be used to think about the age profile of mortality among the oldest old. We have assumed that it is possible to directly manipulate or choose the level of resources I which measured in this model in health units. But individuals cannot in fact directly choose levels of I. Instead, like in the Grossman model, they choose inputs into health: they decide how much food, exercise, medicine, alcohol to consume, based on how much utility they derive from these items directly and on how these affect their health. With data on inputs this would be worth modeling. Lastly the model's implications for wages, consumption and health care expenditures can be improved and taken to data for more recent cohorts. Particularly the model has implications for the evolution of health (and productivity) at the individual level, which can be investigated with panel data. We leave these applications to future work.

References

Alemayehu, Berhanu and Kenneth E. Warner, "The lifetime distribution of health care costs," *Health services research*, 2004, *39* (3), 627–642.

Almond, Douglas and Janet Currie, "Killing me softly: The fetal origins hypothesis," *The Journal of Economic Perspectives*, 2011, *25* (3), 153–172.

_ , _ , and Valentina Duque, "Childhood Circumstances and Adult Outcomes: Act II," Technical Report, National Bureau of Economic Research 2017.

Bardet, Jean-Pierre, K.-A. Lynch, G.-P. Mineau, Michael Hainsworth, and Matthew Skolnick, "La mortalité maternelle autrefois: une étude comparée (de la France de l'ouest à l'Utah)," in "Annales de démographie historique," Vol. 1981 Persée-Portail des revues scientifiques en SHS 1981, pp. 31–48.

Bouvier-Colle, M.-H., F. Pequignot, and E. Jougla, "Mise au point sur la mortalité maternelle en France: fréquence, tendances et causes," 2008.

Bozzoli, Carlos, Angus Deaton, and Climent Quintana-Domeque, "Adult height and childhood disease," *Demography*, 2009, *46* (4), 647–669.

Bronikowski, Anne M., Jeanne Altmann, Diane K. Brockman, Marina Cords, Linda M. Fedigan, Anne Pusey, Tara Stoinski, William F. Morris, Karen B. Strier, and Susan C. Alberts, "Aging in the natural world: comparative data reveal similar mortality patterns across primates," *Science*, 2011, *331* (6022), 1325–1328.

Bruckner, Tim and Ralph Catalano, "The sex ratio and age-specific male mortality: evidence for culling in utero," *American Journal of Human Biology*, 2007, *19* (6), 763–773.

Carnes, Bruce A., S. Jay Olshansky, and Douglas Grahn, "Continuing the search for a law of mortality," *Population and Development review*, 1996, pp. 231–264.

Carriere, JACQUES F., "Parametric models for life tables," *Transactions of the Society of Actuaries*, 1992, *44*, 77–99.

Case, Anne, Darren Lubotsky, and Christina Paxson, "Economic status and health in childhood: The origins of the gradient," *The American Economic Review*, 2002, *92* (5), 1308–1334.

Chetty, Raj, Michael Stepner, Sarah Abraham, Shelby Lin, Benjamin Scuderi, Nicholas Turner, Augustin Bergeron, and David Cutler, "The association between income and life expectancy in the United States, 2001-2014," *Jama*, 2016, *315* (16), 1750–1766.

Contoyannis, Paul, Andrew M. Jones, and Nigel Rice, "The dynamics of health in the British Household Panel Survey," *Journal of Applied Econometrics*, 2004, *19* (4), 473–503.

Cunha, Flavio and James Heckman, "The Technology of Skill Formation," *The American Economic Review*, 2007, *97* (2), 31–47.

Currie, Janet and Mark Stabile, "Socioeconomic Status and Child Health: Why Is the Relationship Stronger for Older Children?," *The American Economic Review*, 2003, *93* (5), 1813–1823.

Cutler, David M., Adriana Lleras-Muney, Tom Vogl, S. Glied, and P. C. Smith, "The Oxford Handbook of Health Economics," 2012.

de La Rochebrochard, Elise, "Age at puberty of girls and boys in France: Measurements from a survey on adolescent sexuality," *Population: An English Selection*, 2000, pp. 51–79.

Deaton, Angus S. and Christina H. Paxson, "The effects of economic and population growth on national saving and inequality," *Demography*, 1997, *34* (1), 97–114.

_ **and** _ , "Aging and inequality in income and health," *The American Economic Review*, 1998, *88* (2), 248–253.

_ **and Christina Paxson**, "Saving, growth, and aging in Taiwan," in "Studies in the Economics of Aging," University of Chicago Press, 1994, pp. 331–362.

Field, Erica, Omar Robles, and Maximo Torero, "Iodine deficiency and schooling attainment in Tanzania," *American Economic Journal: Applied Economics*, 2009, *1* (4), 140–169.

Finch, Caleb E. and Eileen M. Crimmins, "Inflammatory exposure and historical changes in human lifespans," *Science*, 2004, *305* (5691), 1736–1739.

Gompertz, Benjamin, "A Sketch of an Analysis and Notation Applicable to the Estimation of the Value of Life Contigencies," *Philosophical Transactions of the Royal Society of London*, 1820, *110*, 214–332.

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

_ , "A Supplement to Two Papers Published in the Transactions of the Royal Society," On the Science Connected with Human Mortality;" The One Published in 1820, and the Other in 1825," *Philosophical Transactions of the Royal Society of London*, 1862, *152*, 511–559.

_ , "On one uniform law of mortality from birth to extreme old age, and on the law of sickness," *Journal of the Institute of Actuaries and Assurance Magazine*, 1871, *16* (5), 329–344.

Grossman, Michael, "On the concept of health capital and the demand for health," *Journal of Political economy*, 1972, *80* (2), 223–255.

Heligman, Larry and John H. Pollard, "The age pattern of mortality," *Journal of the Institute of Actuaries*, 1980, *107* (01), 49–80.

Honoré, Bo E. and Adriana Lleras-Muney, "Bounds in competing risks models and the war on cancer," *Econometrica*, 2006, *74* (6), 1675–1698.

Hummer, Robert A. and Joseph T. Lariscy, "Educational attainment and adult mortality," in "International handbook of adult mortality," Springer, 2011, pp. 241–261.

Kohler, Iliana V., Samuel Preston, and Laurie Bingaman Lackey, "Comparative mortality levels among selected species of captive animals," *Demographic Research*, 2006, *15*, 413–434.

Li, Ting and James Anderson, "Shaping human mortality patterns through intrinsic and extrinsic vitality processes," *Demographic research*, 2013, *28*, 341–372.

Loudon, Irvine, "The transformation of maternal mortality.," *BMJ: British Medical Journal*, 1992, *305* (6868), 1557.

_ , "Maternal mortality in the past and its relevance to developing countries today," *The American journal of clinical nutrition*, 2000, *72* (1), 241s–246s.

Murphy, Kevin M. and Robert H. Topel, "The value of health and longevity," *Journal of political Economy*, 2006, *114* (5), 871–904.

Occhino, Filippo, Kim Oosterlinck, and Eugene N. White, "How Occupied France Financed Its Own Exploitation in World War II," Technical Report, National Bureau of Economic Research 2006.

Palloni, Alberto and Hiram Beltrán-Sánchez, "Demographic Consequences of Barker Frailty," in "Dynamic Demographic Analysis," Springer, 2016, pp. 147–176.

Preston, Samuel, Patrick Heuveline, and Michel Guillot, "Demography: measuring and modeling population processes," 2000.

Sharrow, David J. and James J. Anderson, "Quantifying Intrinsic and Extrinsic Contributions to Human Longevity: Application of a Two-Process Vitality Model to the Human Mortality Database," *Demography*, 2016, *53* (6), 2105–2119.

Smith, James P., "Healthy bodies and thick wallets: the dual relation between health and economic status," *The journal of economic perspectives: a journal of the American Economic Association*, 1999, *13* (2), 144.

Thiele, Thorvald Nicolai, "On a mathematical formula to express the rate of mortality throughout the whole of life, tested by a series of observations made use of by the Danish Life Insurance Company of 1871," *Journal of the Institute of Actuaries*, 1871, *16* (5), 313–329.

Thomas, Duncan, Elizabeth Frankenberg, Jed Friedman, Jean-Pierre Habicht, Nathan Jones, Christopher McKelvey, Gretel Pelto, Bondan Sikoki, James P. Smith, Cecep Sumantri, and others, "Causal effect of health on labor market outcomes: Evidence from a random assignment iron supplementation intervention," *California Center for Population Research*, 2004.

Vaupel, James W., Kenneth G. Manton, and Eric Stallard, "The impact of heterogeneity in individual frailty on the dynamics of mortality," *Demography*, 1979, *16* (3), 439–454.

Wagstaff, Adam, "The demand for health: some new empirical evidence," *Journal of Health economics*, 1986, *5* (3), 195–233.

Wilcox, Allen J. and IAN T RUSSELL, "Birthweight and perinatal mortality: I. On the frequency distribution of birthweight," *International Journal of Epidemiology*, 1983, *12* (3), 314–318.

Appendix A1: Notes on the empirical method

1. Data

Territory changes. The table below describes the details of the changes in territory that took place in France since 1816.

Year	Territorial Changes
1861	Annexion of *departements* of Savoie and Haute-Savoie, and of *Comte* de Nice
1869	Franco-Prussian war: loss of Alsace-Lorraine
1914-	WWI: East of France, from Nord Pas-de-Calais to Vosges, is occupied by German military.
1919	At the end of WWI, Alsace-Lorraine is re-integrated to French territory
1939	WW2: Loss of Alsace-Lorraine
1943	WW2: Loss of Corsica
1945	Current territory: Alsace-Lorraine and Corsica are re-integrated to French territory

These changes in territory results in large changes in the population and death counts. This is illustrated below for population. It is unclear how to compute mortality in the year of the change. We compute it by using a weighted average of the population at the beginning and end of the year..

Migration. In the HMD cohort population counts are available. However, because of migrations, these counts cannot be used to derive a survival curve for a cohort. Because of net positive immigration occurring in France, the number of individuals in a given cohort can even increase from one year to the next. This is especially true at the end of the Algerian War. (e.g. the size of the female cohort born in 1910 increases from 300,369 to 303,273 between 1962 and 1963, despite a reported mortality rate of 0.5162. . The unit of analysis in our model of mortality is a country cohort, hence abstracts from migration. In our model the mortality rates coincide exactly with the slope of the survival curve. This is not true in the HMD. The population of the cohort melts natives and immigrants of the same age.

2. Computing the death rates, survival rates and life expectancy

Death rates. When taking our model to the data we target the most direct counterpart of our modeled cohort "mortality rate", which is computed as the number of individuals who died during a year, divided by the number of individuals alive at the beginning of the day. In typical life tables this number corresponds to what demographer call q_t, the probability of dying in a given year, and is conceptually distinct to the mortality rate, denoted by m_t. The main difference lies in adjusting the denominator — the size of the population. As more individuals die during the year the population needs to be adjusted to estimate the size of the remaining population exposed to the risk of death. Because our baseline model does not take this adjustment into account, we compute a direct counterpart of our theoretical object. Therefore, we compute the raw death rate in year t for a given cohort , q_t , as follows:

$$q_t = \frac{D_t}{N_t}$$

where D_t is the death count for year t from the HMD cohort table and N_t is the population on January 1st of year t. The HMD makes adjustments to compute a probability that is corrected for the fact that the data do not tract the same individuals over time, so the probability of dying is not correctly computed for a given cohort. The q we estimate with the raw counts is very similar to what is reported by the HMD except

for the first year of life and the last years of life as shown in Figure 26. This results in our under-estimating life expectancy somewhat.

Survival curves. We compute the survival curve recursively as follows. After initializing $S_0 = 100$, we iteratively compute:

$$S_t = S_{t-1} \times (1 - q_{t-1})$$

Life expectancy. Life Expectancy (LE) is an important statistics for the health profile of a given cohort. We compute LE as a way of comparing our model to the data in a parsimonious way. While we try to provide informative estimates of cohort life expectancy, we do not claim that their accuracy is comparable to demographic studies. Nevertheless, as we treat the series generated by our model in exactly the same manner as the data series, we obtain pairs of LE that are readily comparable.

4. Estimation routine

We compute our estimates using Matlab's canned fminsearch routine, a downhill simplex method, and Powell (1964)'s conjugate direction method. We first estimate the model using fminsearch until the objective function changes by less than XXX. We then use these as starting values for Powell's routine. Once Powell's routine converges, we use the estimated values from this procedure and implement fminsearch again until it converges.

5. Bootstrapping standard errors

Estimates from sample data come with standard errors. However, the mortality rates in the HMD are computed from birth certificates of the total population, not a sample of it. A typical cohort in our study counts 400,000 individuals. As a results, the s.e. are extremely small . As in we do not report them for the French cohorts.

In contrast, we do compute the standard errors for the monkey estimates as the data in that case consist of samples of one or two hundreds of individuals.

How would one bootstrap errors? Given a series of mortality rates for a cohort, a sample of size N can be viewed as a sequence of Bernoulli trials with varying success rates.

Alternatively, one can view the survival curve of a population of size N as an $N \times 1$ vector of age at death. One can produce bootstrap estimates by drawing with replacement M subsamples of size S and compute the empirical survival curve.

Appendix A3: Proofs omitted in the text

This is the model that we study for which we will prove propositions 1, 2 and 3 in the text:

$$D_0 = 0,\ H_0 \sim \mathcal{N}\left(\mu_H, \sigma_H^2\right)$$

$$H_t = H_{t-1} - \delta \cdot t^{\alpha} + I + \varepsilon_t \text{ if } D_{t-1} = 0,\ \varepsilon_t \sim \mathcal{N}\left(0, \sigma_\varepsilon^2\right)$$

$$D_t = \mathbb{I}(H_t \leq \underline{H}, D_{t-1} = 0)$$

with $\delta \in (0, \infty), \alpha \in (0, \infty)$, and $I \in \mathbb{R}$. $\underline{H}$ and σ_H^2 are normalized to be 0 and 1, respectively. We now describe the behavior of this model and then go on to analyze the effect of changes in each of its underlying parameters.

Let $\hat{H}_t \equiv \mathbb{E}\left[H_t \mid H_t > 0\right]$ denote the average health in the living population with age t and $\sigma_{\hat{H}_t} \equiv Var\left[H_t \mid H_t > 0\right]$ the variance of health among the living.

Proposition 1: Identification

Below we formally prove that the baseline model is identified.

Formal approach

Suppose we have two sets of parameters $\theta = (I, \delta, \sigma_\varepsilon, \alpha, \mu_H)$ and $\theta' = (I', \delta', \sigma'_\varepsilon, \alpha', \mu'_H)$.

We say that θ and θ' are **observationally equivalent** (OE) if they imply the same mortality rates at each age, i.e. iff

$$MR_t(\theta) = MR_t(\theta'), \forall t \in \mathbb{N}$$

Equivalently, we could define observational equivalence in terms of survival rates $\{S_t(\theta)\}_{t>0}$ since each sequence can be uniquely recovered from the other one.

We say that θ and θ' are **weakly observationally equivalent** (WOE) if and only if 1) θ and θ' are OE and 2) they do not generate the same sequences of health distributions, i.e.

$$\begin{cases} \exists t \in \mathbb{N}, \exists x \in \mathbb{R}^+ & F_{H_t}(x;\theta) \neq F_{H_t}(x;\theta') \\ \forall t \in \mathbb{N} & MR_t(\theta) = MR_t(\theta') \end{cases}$$

We say that θ and θ' are **strongly observationally equivalent** (SOE) iff 1) θ and θ' are OE and 2) they generate the same sequences of health distributions, i.e.

$$\begin{cases} \forall t \in \mathbb{N}, & F_{H_t}(\cdot;\theta) = F_{H_t}(\cdot;\theta') \\ \forall t \in \mathbb{N} & MR_t(\theta) = MR_t(\theta') \end{cases}$$

Although we cannot distinguish between OE and WOE using only mortality rate, we could potential observe some other features of the distributions of health at all ages that could break the identification. A step in that direction would be to observe for a cohort a good proxy for health.

Suppose that θ and θ' are OE, then here are two cases to consider, either they are WOE or SOE. In the following we show successively that neither WOE or SOE is possible.

Now suppose that $(I, \delta, \sigma, \alpha, \mu_H) \neq (I', \delta', \sigma, \alpha', \mu'_H)$

Case 1: SOE

To show that θ and θ' cannot be SOE, let's work towards a contradiction.

Suppose that $(I, \delta, \alpha, \mu_H) \neq (I', \delta', \alpha', \mu'_H)$ then it is clear that the first 4 modes cannot be equal (4 non-linear polynomial equations).

$$\begin{aligned} mode(1) &= \mu_H + I - \delta \\ mode(2) &= \mu_H + 2I - \delta(1 + 2^\alpha) \\ mode(3) &= \mu_H + 3I - \delta(1 + 2^\alpha + 3^\alpha) \\ mode(4) &= \mu_H + 4I - \delta(1 + 2^\alpha + 3^\alpha + 4^\alpha) \end{aligned} \tag{1}$$

Now suppose that $(I, \delta, \alpha, \mu_H) = (I', \delta', \alpha', \mu'_H)$ and $\sigma_\varepsilon \neq \sigma'_\varepsilon$ then the first two mortality rates cannot be equal

$$m_1(\theta) = \Phi\left(\frac{-\mu_H - I + \delta}{\sqrt{1 + \sigma_\varepsilon^2}}\right) \neq \Phi\left(\frac{-\mu_H - I + \delta}{\sqrt{1 + (\sigma'_\varepsilon)^2}}\right) = m_1(\theta') \tag{2}$$

Case 2: WOE

Consider the first case. Let $\tau \equiv \min\{t \in \mathbb{N} \mid \exists x \in \mathbb{R}^+ \, F_{H_t}(x;\theta) \neq F_{H_t}(x;\theta')\}$. It is well defined by definition of WOE.

Notice that because F_{H_t} is continuous except at 0, we can assume wlog that the two cdf differ on some non trivial interval $(a, b) \supset \{x\}$

If $\tau > 1$ then because at the previous period the two distribution are the same, it must be the case that $\sigma_\varepsilon \neq \sigma'_\varepsilon$, but then the

So it must be that $\tau = 1$ i.e. the distribution start differing at the first period.

Some useful lemmas

The variance is separately identified.

If $\{S_t\}_{t\in[0,T]}$ is observed for an arbitrary large T. Then the variance is identified. Intuitively, the variance of the shock characterizes the thickness of the right-hand tail. If one population has a larger variance than the other one then the ratio of survivors grows arbitrarily large at old age.

More formally, let $\theta = \{\sigma, \psi\}$ denote the set of parameters. For any ψ, ψ' such that

$$\sigma > \sigma' \Longrightarrow \lim_{t\to+\infty} \frac{S_t(\sigma, \psi)}{S_t(\sigma', \psi')} = +\infty$$

Proof:

Let $\psi = \{\alpha, \delta, \mu_0, \kappa\}$ and $\sigma > \sigma'$.

We have that for any t $\lim_{t\to+\infty} \frac{f_{H_t}(x;\sigma,\psi)}{f_{H_t}(x;\sigma',\psi')} = +\infty$

The only way to "compensate" for a small variance, which creates in old ages a right tail of very healthy people is to have a lower depreciation (δ and α). However because the tail decreases at exponential rate, we have

$$\lim_{t\to+\infty} \frac{f_{H_t}(x + z_t;\sigma,\psi)}{f_{H_t}(x;\sigma',\psi')} = +\infty \; \forall x > 0$$

where $z_t = \sum_{s<t} \delta t^\alpha - \sum_{s<t} \delta' t^{\alpha'}$

Remark: as is well-known with this kind of "identification at infinity" (see Chamberlain 1986, and Heckman 1990) in practice the

Single-peakedness of f_{H_t}

For any t, one of these cases occurs: either (1) f_{H_t} is hump-shaped (increasing then decreasing) or (2) f_{H_t} is strictly decreasing.

Proof:

We start with a single-peaked distribution.

Now single-peakedness is preserved when we take the convolution with an independently distributed random variable .

At some point the truncation "eats" all the part to the left of the hump.

Corollary: Mode of f_{H_t} There exists $t_{mode} > 0$ such that

$$\begin{cases} mode\left(f_{H_t}\right) > 0 & t < t_{mode} \\ mode\left(f_{H_t}\right) = 0 & t \geq t_{mode} \end{cases}$$

And the mode of the distribution is

$$\max\left\{\mu_0 + I \cdot t - \delta \sum_{s=0}^{t} s^{\alpha}, 0^{+}\right\}$$

Proposition 2: Basic Properties of the model

Basic Properties of the model:

1. Everyone dies with probability 1: $\lim_{t\to\infty} Pr(H_t = 0) = 1$.
2. For sufficiently high I (relative to σ_{ε}^2 and σ_H^2) mortality rates declines (up to age t_1) and then increases with age: $MR_t - MR_{t-1} \leqslant 0$ if $t \leqslant t_1$ and $MR_t - MR_{t-1} \geq 0$ if $t > t_1$.
3. The average health of the living increases and then decreases with age: $\hat{H}_t - \hat{H}_{t-1} \leqslant 0$ if $t \leqslant t_2$ and $\hat{H}_t - \hat{H}_{t-1} \geq 0$ if $t > t_2$.
4. The variance of health among the living increases and then falls: $\sigma_{\hat{H}_t} - \sigma_{\hat{H}_{t-1}} \leq 0$ if $t \leqslant t_3$ and $\sigma_{\hat{H}_t} - \sigma_{\hat{H}_{t-1}} \geq 0$ if $t > t_3$.

1. **Everyone dies eventually.**

   Consider the process $\{H_t^*\}_{t=1}^{\infty}$, defined by $H_0^* = H_0 \sim \mathcal{N}\left(\mu_H, \sigma_H^2\right)$ and the recurrence relation:

   $$H_t^* = H_{t-1}^* + I - \delta \cdot t^{\alpha} + \varepsilon_t \ , \ \varepsilon_t \sim \mathcal{N}\left(0, \sigma_{\varepsilon}^2\right) \tag{3}$$

   It is easy to tell that $0 \leq P\left(H_t > z\right) \leq P\left(H_t^* > z\right)$ for any $z > 0$.

   Now for any $t \geq 0$, H_t^* is normally distributed with mean

   $$\mu_{H_t^*} = \mu_H + I \cdot t - \delta \sum_{k=1}^{t} k^{\alpha} \tag{4}$$

and standard deviation

$$\sigma_{H_t^*} = \sqrt{\sigma_H^2 + t \cdot \sigma_\varepsilon^2} \tag{5}$$

Hence, $P\left(H_t^* > z\right) = 1 - \Phi(\frac{z - \mu_{H_t^*}}{\sigma_{H_t^*}})$, where Φ is the CDF of the standard normal distribution.

As $t \to \infty$, we have $\mu_{H_t^*} \sim I \cdot t - \delta \cdot \frac{t^{\alpha+1}}{\alpha+1}$ and $\sigma_{H_t^*} \sim \sqrt{t} \cdot \sigma_\varepsilon$.

Therefore if $\alpha > 0$, $\frac{\mu_{H_t^*}}{\sigma_{H_t^*}} \to -\infty$ as $t \to \infty$.

2. **The mortality rates are U-shaped under suitable parametric restrictions.**

   The mortality rate at time t is defined as

$$MR_t = \frac{F_{H_t}(0) - F_{H_{t-1}}(0)}{1 - F_{H_{t-1}}(0)} \tag{6}$$

   First, as long as σ_ε small enough w.r.t 1 , then $MR_1 > MR_2$, then again we need δ, α to be small enough relative to I and σ_ε so that the aging is not too strong, otherwise MR increases immediately at age 2.

   The sufficient condition for

$$MR_t - MR_{t-1} \leqslant 0 \tag{7}$$

   is

$$F_{H_t}(0) + F_{H_{t-2}}(0) - 2F_{H_{t-1}} + F_{H_{t-1}}^2(0) - F_{H_t}(0)F_{H_{t-2}}(0) \leqslant 0 \tag{8}$$

3. **Average health has inverted U-shape.**

Proposition 3: Comparative statics

1. Increasing the investment I or the average health at birth μ_H unambiguously decreases mortality at all ages: $\frac{\partial MR_t}{\partial I} \leq 0$, $\frac{\partial MR_t}{\partial \mu_H} \leq 0$.[39]

2. Increasing any of the aging parameters, δ or α, unambiguously increases mortality at all ages: $\frac{\partial MR_t}{\partial \delta} \geq 0, \frac{\partial MR_t}{\partial \alpha} \geq 0$.

3. An increase in σ_H^2 can increase or decrease the mortality rate at a given age. An increase in σ_H^2 increases the mortality rate at young ages $\frac{\partial MR_t}{\partial \sigma_H^2} \geq 0$ if $\delta t^\alpha \leq I$. Ultimately, an increase in σ_H^2 generates selection and reduces mortality in the very old age, for some t_σ, $\frac{\partial MR_{t+s}}{\partial \sigma_H^2} < 0, \forall$s>$t_\sigma$.

4. Investment and health at birth are complements: $\frac{\partial^2 MR_t}{\partial I \partial \mu_H} \leq 0$.

   1. **Increasing the investment** I:

   In period 1, we have $\forall z \geqslant 0$,

$$\begin{aligned} F_{H_1}(z) &= Pr(H_0 + I - \delta + \varepsilon_1 \leqslant z) \\ &= Pr(H_0 + \varepsilon_1 \leqslant z + \delta - I) \\ &= \Phi\left(\frac{z + \delta - I - \mu_H}{\sqrt{1 + \sigma_\varepsilon^2}}\right) \end{aligned} \tag{9}$$

[39] Changing the threshold also affects mortality rates negatively throughout the lifetime.

$$
\begin{aligned}
MR_1 &= Pr(H_1 \leqslant 0) \\
&= Pr(H_0 + I - \delta + \varepsilon_1 \leqslant 0) \\
&= \Phi\left(\frac{\delta - I - \mu_H}{\sqrt{1+\sigma_\varepsilon^2}}\right)
\end{aligned} \tag{10}
$$

It is easy to see that as I or μ_H increases, both $F_{H_1}(z)$ and MR_1 decreases.

$\forall t \geqslant 2, \forall z \geqslant 0$, We have

$$
\begin{aligned}
F_{H_t}(z) &= Pr(H_t \leqslant z) \\
&= Pr(H_t \leqslant z, H_{t-1} > 0) + Pr(H_t \leqslant z, H_{t-1} \leqslant 0) \\
&= Pr(H_{t-1} + I - \delta t^\alpha + \varepsilon_t \leqslant z, H_{t-1} > 0) + Pr(H_{t-1} \leqslant 0) \\
&= Pr(0 < H_{t-1} \leqslant z + \delta t^\alpha - I - \varepsilon_t) + Pr(H_{t-1} \leqslant 0) \\
&= \frac{1}{\sigma_\varepsilon}\int_{x=0}^{\infty} \phi\left(\frac{x - z - \delta t^\alpha + I}{\sigma_\varepsilon}\right)\left(F_{H_{t-1}}(x) - F_{H_{t-1}}(0)\right)dx + F_{H_{t-1}}(0)
\end{aligned} \tag{11}
$$

$$
MR_t = \frac{F_{H_t}(0) - F_{H_{t-1}}(0)}{1 - F_{H_{t-1}}(0)} \tag{12}
$$

Then

$$
\begin{aligned}
&\frac{\partial F_{H_t}(z;I)}{\partial I} \\
=&\frac{\partial}{\partial I}\left(\frac{1}{\sigma_\varepsilon}\int_{x=0}^{\infty} \phi\left(\frac{x - z - \delta t^\alpha + I}{\sigma_\varepsilon}\right)\left(F_{H_{t-1}}(x) - F_{H_{t-1}}(0)\right)dx\right) + \frac{\partial F_{H_{t-1}}(0;I)}{\partial I} \\
=&\frac{1}{\sigma_\varepsilon^2}\int_{x=0}^{\infty} \phi'\left(\frac{x - z - \delta t^\alpha + I}{\sigma_\varepsilon}\right)\left(F_{H_{t-1}}(x) - F_{H_{t-1}}(0)\right)dx \\
&+ \frac{1}{\sigma_\varepsilon}\int_{x=0}^{\infty} \phi\left(\frac{x - z - \delta t^\alpha + I}{\sigma_\varepsilon}\right)\left(\frac{\partial F_{H_{t-1}}(x;I)}{\partial I} - \frac{\partial F_{H_{t-1}}(0;I)}{\partial I}\right)dx \\
&+ \frac{\partial F_{H_{t-1}}(0;I)}{\partial I} \\
=&\frac{1}{\sigma_\varepsilon^2}\int_{x=0}^{\infty} \phi'\left(\frac{x - z - \delta t^\alpha + I}{\sigma_\varepsilon}\right)\left(F_{H_{t-1}}(x) - F_{H_{t-1}}(0)\right)dx \\
&+ \frac{1}{\sigma_\varepsilon}\int_{x=0}^{\infty} \phi\left(\frac{x - z - \delta t^\alpha + I}{\sigma_\varepsilon}\right)\frac{\partial F_{H_{t-1}}(x;I)}{\partial I}dx \\
&+ \Phi\left(-z - \delta t^\alpha + I\right)\frac{\partial F_{H_{t-1}}(0;I)}{\partial I}
\end{aligned} \tag{13}
$$

All three items are negative, so we have

$$
\frac{\partial F_{H_t}(z;I)}{\partial I} \leqslant 0 \tag{14}
$$

$$\begin{aligned}&\frac{\partial MR_t}{\partial I}\\ =&\frac{1}{(1-F_{t-1}(0))^2}\left(\frac{\partial F_{H_t}(0)}{\partial I}(1-F_{H_{t-1}}(0))-\frac{\partial F_{H_{t-1}}(0)}{\partial I}(1-F_{H_t}(0))\right)\end{aligned} \tag{15}$$

2. Increasing any of the aging parameters, δ or α

Let $a_t = I - \delta t^\alpha$. The random variable H_t has a mass point at $z = 0$ but is continuous on $(0, +\infty)$. $F_{H_t}(0)$ is the probability of not surviving until age t while for any $z > 0$, the cdf can be expressed

$$F_{H_t}(z) = \int_{x=0}^{\infty} \Phi\left(\frac{z - x - a_t}{\sigma_\varepsilon}\right) f_{H_{t-1}}(x)\, dx + F_{H_{t-1}}(0)$$

Equivalently, after integration by parts, one obtains:

$$F_{H_t}(z) = -\frac{1}{\sigma_\varepsilon}\int_{x=0}^{\infty} \phi\left(\frac{z - x - a_t}{\sigma_\varepsilon}\right) F_{H_{t-1}}(x)\, dx + F_{H_{t-1}}(0)$$

Hence the mortality rate at age t, which is the probability of dying at age t conditional on surviving until age t, can be written:

$$MR_t = \frac{F_{H_t}(0) - F_{H_{t-1}}(0)}{1 - F_{H_{t-1}}(0)}$$

Suppose that for every t we increase the constant investment level I to some level $I' > I$. Following the expression above, the impact can be be decomposed in two: first, a direct effect on the probability of dying at age t (the numerator) and, second, a compounded effect carried through the distribution of health for those attaining age t.We show that, for any t, both effects go in the same direction: an increase in I simultaneously increases the probability of surviving untill age t (hence increases the denominator) and reduces the probability of dying at age t (the numerator goes down). We prove the following lemma.

For all t, we have:

1. $\forall z \geqslant 0$, $\frac{\partial F_{H_t}(z;I)}{\partial I} \leq 0$
2. $\frac{\partial MR_t}{\partial I} \leq 0$

Note: I don't think we can prove Lemma 1. The proof here is wrong.

We prove these inequalities jointly and by induction.

Notice that $\frac{\partial F_{H_t}(\cdot;I)}{\partial I} \leq 0$ signifies that the cdf's are ranked by first order stochastic dominance. The higher the I, the further the distribution is pushed to the right, which decreases the value of the cdf at any point x as I increases. Because all the individuals would then be in better health, ceteris paribus, fewer of them will die each period. Combined with a higher denominator, this delivers a lower mortality rate at each point.

At $t = 0$: $F_{H_1}(z;I) = \Phi\left(\frac{z-\mu_0}{\sigma_0}\right)$ hence $\frac{\partial F_{H_1}(z;I)}{\partial I} = 0$. $MR_t = F_{H_t}(0) = \Phi\left(\frac{z-\mu_0}{\sigma_0}\right)$ which, again, is non-increasing with I.

For any $t \geq 1$, suppose that $\frac{\partial F_{H_{t-1}}}{\partial I} \leq 0$ and $\frac{\partial MR_{t-1}}{\partial I} \leq 0$.

Let's first focus on the first claim:

$$\frac{\partial F_{H_t}(z;I)}{\partial I} = \frac{\partial}{\partial I}\left[-\frac{1}{\sigma_\varepsilon}\int_{x=0}^{\infty} \phi\left(\frac{z - x - a_t}{\sigma_\varepsilon}\right) F_{H_{t-1}}(x)\, dx\right] + \frac{\partial F_{H_{t-1}}(0;I)}{\partial I}$$

The second term is negative, by assumption, while the first term is equal to

$$\frac{1}{\sigma_\varepsilon^2}\int_{x=0}^{\infty}\phi'\left(\frac{z-x-a_t}{\sigma_\varepsilon}\right)F_{H_{t-1}}(x)\,dx-\frac{1}{\sigma_\varepsilon}\int_{x=0}^{\infty}\phi\left(\frac{z-x-a_t}{\sigma_\varepsilon}\right)\frac{\partial F_{H_{t-1}}(z;\,I)}{\partial I}dx$$

Again, by assumption, $\frac{\partial F_{H_{t-1}}(z;I)}{\partial I}\leq 0$, which takes care of the rightmost term.

Now, consider the change of variable $u=\frac{z-x-a_t}{\sigma_\varepsilon}$. We can rewrite the leftmost term:

$$-\frac{1}{\sigma_\varepsilon^3}\int_{u=-\infty}^{\frac{z-a_t}{\sigma_\varepsilon}}\phi'(u)\,F_{H_{t-1}}(z-a_t-\sigma_\epsilon u)\,du$$

There are two cases. If $\frac{z-a_t}{\sigma_\varepsilon}\leq 0$ then the integrand is always positive as $\phi'>0$ for negative real numbers, and we conclude that $\frac{\partial F_{H_t}(z;I)}{\partial I}\leq 0$. If $\frac{z-a_t}{\sigma_\varepsilon}>0$ then we can split the integral in three terms:

$$\begin{aligned}
&-\tfrac{1}{\sigma_\varepsilon^3}\textstyle\int_{u=-\infty}^{-\frac{z-a_t}{\sigma_\epsilon}}\phi'(u)\,F_{H_{t-1}}(z-a_t-\sigma_\varepsilon u)\,du\\
&-\tfrac{1}{\sigma_\varepsilon^3}\textstyle\int_{u=-\frac{z-a_t}{\sigma_\varepsilon}}^{0}\phi'(u)\,F_{H_{t-1}}(z-a_t-\sigma_\varepsilon u)\,du\\
&-\tfrac{1}{\sigma_\varepsilon^3}\textstyle\int_{u=0}^{\frac{z-a_t}{\sigma_\varepsilon}}\phi'(u)\,F_{H_{t-1}}(z-a_t-\sigma_\varepsilon u)\,du\\
=\quad&-\tfrac{1}{\sigma_\varepsilon^3}\int_{u=-\infty}^{-\frac{z-a_t}{\sigma_\varepsilon}}\phi'(u)\,F_{H_{t-1}}(z-a_t-\sigma_\varepsilon u)\,du\\
-\quad&\tfrac{1}{\sigma_\varepsilon^3}\int_{u=-\frac{z-a_t}{\sigma_\varepsilon}}^{0}\phi'(u)\left[F_{H_{t-1}}(z-a_t-\sigma_\varepsilon u)-F_{H_{t-1}}(z-a_t+\sigma_\varepsilon u)\right]du
\end{aligned}$$

(as $\phi'(-u)=-\phi'(u)$ and the cdf $F_{H_{t-1}}$ is non-decreasing). This proves that $\frac{\partial F_{H_t}(z;I)}{\partial I}\leq 0$ for any $z\in\mathbb{R}$.

With that result in hand, it is easy to prove that $\frac{\partial MR_t}{\partial I}\leq 0$. Setting $z=0$, it follows directly that the denominator decreases with I. Regarding the numerator we have

$$\begin{aligned}
\tfrac{\partial}{\partial I}\left[F_{H_t}(0;\,I)-F_{H_{t-1}}(0;\,I)\right]&=\tfrac{\partial F_{H_t}(0;\,I)}{\partial I}-\tfrac{\partial F_{H_{t-1}}(0;\,I)}{\partial I}\\
&=\tfrac{\partial}{\partial I}\left[-\tfrac{1}{\sigma_\varepsilon}\int_{x=0}^{\infty}\phi\left(\tfrac{x-a_t}{\sigma_\varepsilon}\right)F_{H_{t-1}}(x)\,dx\right]\\
&\leq 0
\end{aligned}$$

since this is the same integral analyzed at the previous step, with $z=0$.

By induction, i) and ii) hold for any $t\geq 1$.

For μ_H same proof, except effect on first period distribution. The successive cdf's inherit the first order stochastic dominance property.

3. An increase in σ_H^2.

The exact same proof applies for δ and α as their impact on F_{H_t} through the aging function a_t is similar to the effect of I.

It can be seen right away that this proof will not work for σ_ε nor σ_0. Increasing any of these variances, - a mean-preserving spread - will not give rise to the first order stochastic ranking of the cdf's that we have used.

Increasing MR on impact. The numerator of the MR_t is given by

$$F_{H_t}(0;\,I)-F_{H_{t-1}}(0;\,I)=\int_{x=0}^{\infty}\Phi\left(\frac{\delta t^\alpha-I-x}{\sigma_\varepsilon}\right)f_{H_{t-1}}(x)\,dx$$

Since Φ is nondecreasing, if one decreases σ_ε at time t, and at this period only, then this expression is necessarily decreasing in σ_ε if $\delta t^\alpha\leq I$.

This follows from the fact that a higher $\sigma_{\varepsilon,t}$ will generate a fatter right-hand tail. For instance, $\lim_{x\to+\infty}\frac{f_{H_t}(x;\sigma_\varepsilon)}{f_{H_t}(x;\sigma'_\varepsilon)}=0$. Now if $\sigma_{\varepsilon,t}$ is changed only at period . From then on, the distributions are modified through a similar process. It can be shown that the fatter right-hand tail property will be preserved. In the very old age, only the popupulation in the right-tail have survived, hence the result.

To prove that the fatter right-hand tail property is preserved, proceed similarly by inference.

Remark 1 Lemma 1 is actually a subcase of the following result, which is slightly stronger:

Suppose that the level of investment is allowed to change at every period, and denote $\mathcal{I}=\{I_1,I_2,...,\}$ and $\mathcal{I}'=\{I'_1,I'_2,...,\}$ two investment sequences. The following holds:

$$\forall s\geq 1,\ I'_s\geq I_s\Longrightarrow\forall t\geq 1,\forall z>0,F_{H_t}\left(z;\mathcal{I}\right)\leq F_{H_t}\left(z;\mathcal{I}'\right)\text{ and }MR_t\left(\mathcal{I}\right)\leq MR_t\left(\mathcal{I}'\right)$$

The mechanics of the proof is almost exactly similar. Increasing investment at any period generates a persistent relation of first-order stochastically dominance in the CDF of health.

4. Investment and health at birth are complements

The proof here is wrong.

going back to the proof of Proposition 1

$$\begin{array}{rcl}\frac{\partial^2 F_{H_{t_2}}(z;\mathcal{I})}{\partial I_{t_1}\partial I_{t_2}} & = & \frac{\partial}{\partial I_{t_1}}\frac{\partial}{\partial I_{t_2}}\left[-\frac{1}{\sigma_\varepsilon}\int_{x=0}^{\infty}\phi\left(\frac{z-x-I_{t_2}+\delta(t_2)^\alpha}{\sigma_\varepsilon}\right)F_{H_{t_2-1}}\left(x,\mathcal{I}\right)dx\right]\\ & + & \frac{\partial}{\partial I_{t_1}}\frac{\partial F_{H_{t_2-1}}(0;\mathcal{I})}{\partial I_2}\\ & = & \frac{\partial}{\partial I_{t_1}}\left[\frac{1}{\sigma_\varepsilon^2}\int_{x=0}^{\infty}\phi'\left(\frac{z-x-a_t}{\sigma_\varepsilon}\right)F_{H_{t_2-1}}\left(x,\mathcal{I}\right)dx\right]+0\\ & = & \frac{1}{\sigma_\varepsilon^2}\int_{x=0}^{\infty}\phi'\left(\frac{z-x-a_t}{\sigma_\varepsilon}\right)\frac{\partial}{\partial I_{t_1}}F_{H_{t_2-1}}\left(x,\mathcal{I}\right)dx\\ & \leq & 0\end{array}$$

$$content... \tag{16}$$

because $\frac{\partial}{\partial I_{t_1}}F_{H_{t_2-1}}\left(x,\mathcal{I}\right)\leq 0$ (increasing investment at time 1 creates a FOSD distribution)

And as a consequence the denominator $1-F_{H_{t_2-1}}\left(0\right)$ goes up as well.

The correct equation should be

$$\begin{array}{l}\frac{\partial^2 F_{H_{t_2}}(z)}{\partial I_{t_1}\partial I_{t_2}}\\ =\frac{\partial}{\partial I_{t_1}\partial I_{t_2}}\left(\frac{1}{\sigma_\varepsilon}\int_{x=0}^{\infty}\phi\left(\frac{x-z-\delta t_2^\alpha+I_{t_2}}{\sigma_\varepsilon}\right)\left(F_{H_{t_2-1}}(x)-F_{H_{t_2-1}}(0)\right)dx\right)+\frac{\partial^2 F_{H_{t_2-1}}(0)}{\partial I_{t_1}\partial I_{t_2}}\\ =\frac{\partial}{\partial I_{t_1}}\left(\frac{1}{\sigma_\varepsilon^2}\int_{x=0}^{\infty}\phi'\left(\frac{x-z-\delta t_2^\alpha+I_{t_2}}{\sigma_\varepsilon}\right)\left(F_{H_{t_2-1}}(x)-F_{H_{t_2-1}}(0)\right)dx\right)\\ =\frac{1}{\sigma_\varepsilon^2}\int_{x=0}^{\infty}\phi'\left(\frac{x-z-\delta t_2^\alpha+I_{t_2}}{\sigma_\varepsilon}\right)\left(\frac{\partial F_{H_{t_2-1}}(x)}{\partial I_{t_1}}-\frac{\partial F_{H_{t_2-1}}(0)}{\partial I_{t_1}}\right)dx\end{array} \tag{17}$$

μ_0 and I are complement

Exactly the same as in Proposition 2 as a change in μ_0 is observationally equivalent to a change in I_1.

Appendix Tables and Figures

Table 1: Estimated parameters for female chimpanzees living in the wild

Gender		Basic model	κ_a
Initial mean health	μ_H	0.9783	1.0043
Investment (annual)	I	0.3295	0.3390
Standard Deviation of Shock	σ_e	1.0871	1.1304
Depreciation	δ	0.0560	0.0553
Aging	α	0.7677	0.7820
Adolescent Hump*	κ_a		0.00001
# of individuals at birth		80	80
# of moments reported		55	55
Fit (survival curve)^b		112.50	111.29
Fit (log of q_x)		2.11	2.11
Actual Life Expectancy		15.38$(13.4)^a$	
Predicted Life Expectancy		15.35	15.35

Data sources: Life tables for primates in the wild come from Bronikowski et al. (2011). In the wild population data come from Brazil, Costa Rica, Kenya, Tanzania, Madagascar and Rwanda.
a. Life expectancy in parenthesis corresponds to the one reported in Bronikowski et al. (2011).
b. We target the survival curve and compute the sum of squared errors – the data provided are in the form of survival rates.
*Adolescence starts at age 8.

Table 2: Modeling prime-age mortality. French Women born in 1816

		(0)	(1)	(2)	(3)	(4)
Model for hump: change in...		Baseline	I	$\underline{H}$	σ_e	κ_a
Initial mean health	μ_H	0.9115	0.9115	0.8151	0.7723	0.8634
Investment	I	0.1336	0.1336	0.1315	0.2159	0.4075
Standard Deviation of Shock	σ_e	0.5556	0.5556	0.4830	0.6300	1.0241
Depreciation	δ	0.0010	0.0010	0.0008	0.0009	0.0006
Aging	α	1.4350	1.4350	1.4462	1.5605	1.7849
Adolescent Hump*			0.1336	0.5586	0.9024	0.0086
Fit (survival curve)^		155.06	155.06	123.03	97.04	12.36
Fit (log of q_x)		3.01	3.01	2.87	2.23	0.74
Fit (death distribution)**		6.21	6.21	4.95	2.95	3.35
Actual Life Expectancy				38.25		
Predicted Life Expectancy		38.43	38.43	38.38	38.45	38.28
Counterfactual Life expectancy^^			38.43	40.92	40.38	45.86

*The estimate in this row corresponds to the value of the parameters after the onset of adolescence. Adolescence starts at age = (- 0.0175 x calendar year) + 47.4 for all women, based on the estimates provided in de La Rochebrochard (2000).

**To make the fit of the age distribution comparable across columns we use the (normalized) number of deaths as weights.

^Our main fit criteria is the sum of squared errors of the survival rate at each age We also report the fit as the sum of squared errors of the log of q_x (the probability of dying between ages x and $x+1$) and the distribution of deaths. We don't target these moments directly–we target the survival curve.

^^Counterfactual Life Expectancy is computed by holding all estimated parameters fixed and setting the adolescent hump to 0.

Table 3: Modeling prime-age mortality French Women born in 1860

		(0)	(1)	(2)	(3)	(4)
Model for hump: change in...		Baseline	I	$\underline{H}$	σ_e	κ_a
Initial mean health	μ_H	1.0981	1.0740	1.0748	0.9589	0.9323
Investment	I	0.1501	0.1563	0.1649	0.4879	0.3318
Standard Deviation of Shock	σ_e	0.5916	0.5873	0.5907	1.0471	0.7932
Depreciation	δ	0.0006	0.0005	0.0006	0.0001	0.0004
Aging	α	1.5742	1.5774	1.5889	2.2001	1.7780
Adolescent Hump*	κ_a		0.1380	0.4902	2.1746	0.0071
Fit (survival curve)^		205.54	197.74	177.41	46.34	11.93
Fit (log of q_x)		2.59	2.49	2.57	1.62	0.70
Fit (death distribution)**		7.49	19.78	23.43	9.07	16.48
Actual Life Expectancy				43.80		
Predicted Life Expectancy		43.95	44.04	43.97	43.88	43.85
Counterfactual Life expectancy^^			46.30	45.91	48.96	51.65

*The estimate in this row corresponds to the value of the parameter after the onset of adolescence. Adolescence starts at age = (- 0.0175 x calendar year) + 47.4 for all women except for the 5th column where the timing of adolescence is estimated as following a normal distribution with mean value (- 0.0175 x calendar year) + 47.4, and standard deviation 1.3285 (calculated from the table of 1975 girls) based on the estimates provided in de La Rochebrochard (2000).

**To make the fit of the age distribution comparable across columns we use the (normalized) number of deaths as weights.

^Our main fit criteria is the sum of squared errors of the survival rate at each age We also report the fit as the sum of squared errors of the log of q_x (the probability of dying between ages x and $x+1$) and the distribution of deaths. We don't target these moments directly–we target the survival curve.

^^Counterfactual Life Expectancy is computed by holding all estimated parameters fixed and setting the adolescent hump to 0.

Table 4: Robustness checks for 1816

		(1)	(2)	(3)	(4)	(5)	(6)	(7)
		Basic	κ_b	κ_b at T~$N()$	T~$N()$ estimated	Weight	Target death	Truncation at 90
Initial mean health	μ_H	0.8634	0.8917	0.8635	0.7558	0.8020	0.7327	0.8784
Investment	I	0.4075	0.4322	0.4149	0.3922	0.4180	0.4743	0.4200
Standard Deviation of Shock	σ_e	1.0241	1.0713	1.0367	0.8958	0.9971	0.9930	1.0552
Depreciation	δ	0.0006	0.0005	0.0005	0.0005	0.0006	0.0006	0.0006
Aging	α	1.7849	1.8321	1.8153	1.7910	1.7807	1.7950	1.7973
Adolescent Hump*	κ_a	0.0086	0.0089	0.0089	0.0102	0.0093	0.0108	0.0087
Accident rate before adolescence	κ_b		0.0005					
Mean*				15.6	14.4			
Standard deviation*				1.32	10.65			
Fit (survival curve)^		12.36	13.03	11.49	4.25	15.86	173.74	12.36
Fit (log of q_x)		0.74	0.65	0.57	0.37	0.92	1.93	0.56
Fit (death distribution)**		3.35	39.55	21.82	2.99	6.89	4.48	28.17
Actual Life Expectancy					38.25			
Predicted Life Expectancy		38.28	38.29	38.28	38.27	38.33	39.27	38.27
Counterfactual Life expectancy^^		45.86	46.45	46.08	47.90	46.66	49.49	45.86

*Adolescence starts at age = (- 0.0175 x calendar year) + 47.4 in columns 1, 4 and 5. In column 2 the timing of adolescence is assumed to follow a normal distribution with mean value (- 0.0175 x calendar year) + 47.4, and standard deviation 1.3285, calculated from the table of 1975 girls in de La Rochebrochard (2000). In column 3 we estimate the mean and the standard deviation of the onset of adolescence.

^Our main fit criteria is the sum of squared errors of the survival rate at each age. We also report the fit as the sum of squared errors of the log of q_x (the probability of dying between ages x and $x+1$) and the distribution of deaths. We don't target these moments directly–we target the survival curve.

^^Counterfactual Life Expectancy is computed by holding all estimated parameters fixed and setting the adolescent hump to 0.

**To make the fit of the age distribution comparable across columns we use the (normalized) number of deaths as weights.

In column 4 we target the survival curve but use the (normalized) number of deaths as weights

In column 5 we target the distribution of the age at death and we use the number of deaths as weights.

Table 5: Estimated parameters for WWII for French Women born in 1921

		(1)	(2)	(3)	(4)
Model for WWII: change in...		I	κ_a	σ_e	$\underline{H}$
Initial condition	μ_H	0.9790	1.0837	1.0638	1.0522
Investment	I	0.2985	0.2739	0.2650	0.2385
Standard Deviation of Shock	σ_e	0.4255	0.4561	0.4358	0.3891
Depreciation	δ	0.0007	0.0008	0.0009	0.0008
Aging	α	1.5358	1.5272	1.4961	1.4785
Adolescence Hump*	κ_a	0.0026	0.0030	0.0032	0.0031
WWII Shock**		-0.1173	0.0036	0.3495	0.8919
Fit (survival curve)^		40.62	37.02	38.49	38.53
Fit (log of q_x)		4.87	2.69	3.11	3.29
Fit during WWII (log of q_x)		0.21	0.53	0.68	0.73
% Difference in # deaths during WWII~~		-0.21	-0.36	-0.45	-0.32
Actual Life Expectancy			66.00		
Predicted Life Expectancy		66.03	66.03	66.02	66.03
Counterfactual Life expectancy^^		70.90	66.24	65.94	66.23
Actual Life expectancy in 1946			55.93		
Life expectancy in 1946		55.27	55.25	55.12	55.28
Counterfactual LE in 1946		60.36	55.24	55.05	55.17

*Hump is modeled as a accident rate that starts in adolescence, set to happen at (- 0.0175 * calendar year) + 47.4 based on the estimates provided in de La Rochebrochard (2000).

**The estimates in this row corresponds to the value of the parameter during the war. For example the first column shows that I was about 0.299 throughout life but decreased to -0.117 during the war. The same applies to columns (3) and (4), the standard deviation decreases from 0.436 to 0.350 and the threshold moves from 0 to 0.892. In column 2, we estimate the value of an additional random shock during the war, an approximate 41% decrease relative to the adolescent hump (but since the shock is independent this is only approximate).

^Our main fit criteria is the sum of squared errors of the survival rate at each age We also report the fit as the sum of squared errors of the log of q_x (the probability of dying between ages x and $x+1$). We don't target these moments directly–we target the survival curve.

^^Counterfactual Life Expectancy is computed by holding all estimated parameters fixed and setting the war parameters to 0.

~This is computed as sum of squared errors during the war years. A lower number is better.

~~This is computed as (predicted - actual)/actual

To make the fit of the age distribution comparable across columns we use the (normalized) number of deaths as weights.

Table 6: Estimated parameters for 1919 Flu pandemic for French Women born in 1900

Model pandemic: change in...		(1) I	(2) κ_a	(3) σ_e	(4) $\underline{H}$
Initial mean health	μ_H	0.7718	0.7032	0.6946	0.7322
Investment	I	0.3460	0.3464	0.3548	0.3569
Standard Deviation of Shock	σ_e	0.5976	0.5485	0.5653	0.5863
Depreciation	δ	0.0006	0.0007	0.0006	0.0006
Aging	α	1.6462	1.6143	1.6581	1.6656
Adolescence Hump*	κ_a	0.0049	0.0049	0.0053	0.0050
WW2 (I value during war)	I_w	0.3245	0.0067	0.0048	0.0091
1918 Flu Shock**		-1.1242	0.0225	0.0000	2.2129
Fit (survival curve)		8.49	11.95	15.91	8.65
Fit (log of q_x)		1.27	2.74	2.57	2.06
Fit during WWII (log of q_x)		0.08	0.04	0.04	0.05
% Difference in # deaths during WWII		-0.13	-0.07	-0.01	-0.06
Fit during Epidemics (log of q_x)		0.00	0.11	0.08	0.07
% Difference in # deaths during Epidemics		0.11	1.14	-0.48	0.82
Actual Life Expectancy			53.81		
Predicted Life Expectancy		53.86	53.86	53.84	53.85
Counterfactual: no Flu		55.83	54.55	53.82	54.25
Counterfactual: no WW2		54.02	56.42	56.30	56.29
Counterfactual: no shocks^^		55.98	57.16	56.28	56.75
Actual life expectancy in 1919			53.09		
Life expectancy in 1919		53.04	53.40	52.90	53.34
Counterfactual life expectancy in 1919		55.36	53.40	52.87	53.14

*Hump is modeled as a accident rate that starts in adolescence, set to happen at (-0.0175 x calendar year) + 47.4, based on the estimates provided in de La Rochebrochard (2000).

**We model the war as a temporary change in I. The values reported in this row correspond to the level of I during the war. For example in column (1) I during the war falls from about 0.38 to about 0.045.

***The estimates in this row corresponds to the value of the parameter during the epidemics. For example column (1) shows that I was about 0.38 throughout life but decreased to about -0.92 during an epidemic. The same applies to columns (3) to (4). In column 2, we estimate the value of an additional random shock during an epidemic, in this case it corresponds to an approximate 100% increase relative to the adolescent hump (but since the shock is independent this is only approximate).

^Our main fit criteria is the sum of squared errors of the survival rate at each age We also report the fit as the sum of squared errors of the log of q_x (the probability of dying between ages x and $x+1$). We don't target these moments directly–we target the survival curve.

^^Counterfactual Life Expectancy is computed by holding all estimated parameters fixed and setting the war and pandemic parameters to 0.

**To make the fit of the age distribution comparable across columns we use the (normalized) number of deaths as weights.

Figure 18: Life expectancy, French Women 1816-1970

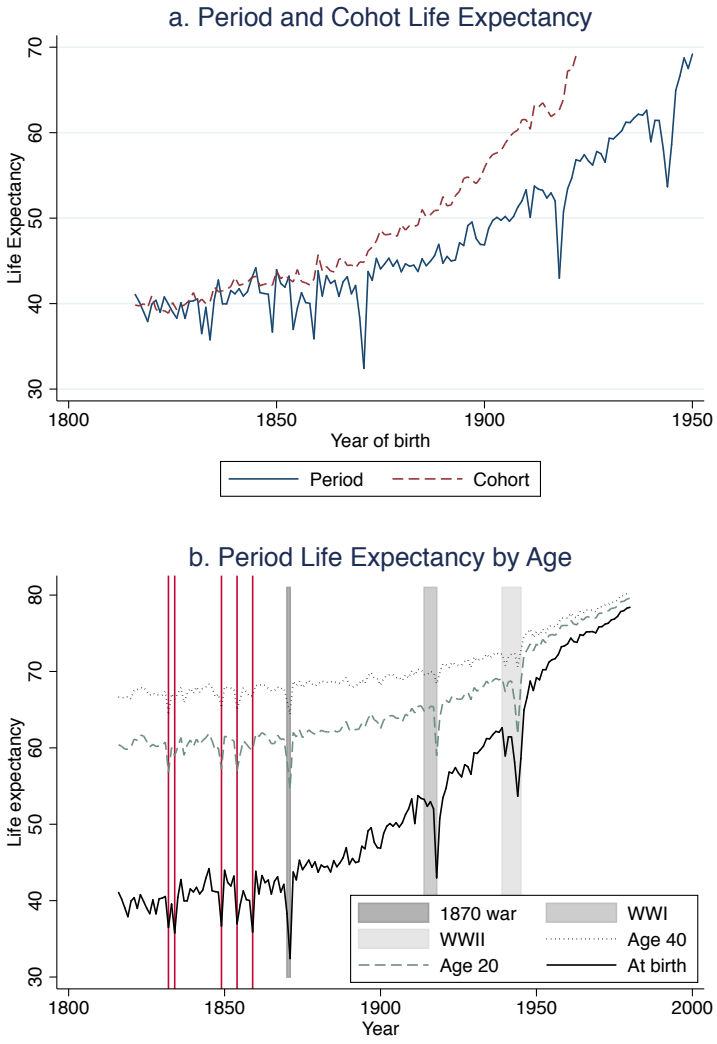

Note: Red lines show 19th century epidemics: cholera (1832, 1849, 1854), influenza (1834), and dysentery (1859). The figure show LE at age 20 + 20, and LE at 40 + 40, something like expectact age at death, rather than remaining years of life left.

Figure 19: Age profile of mortality of women born in France between 1860 and 1940, by decade

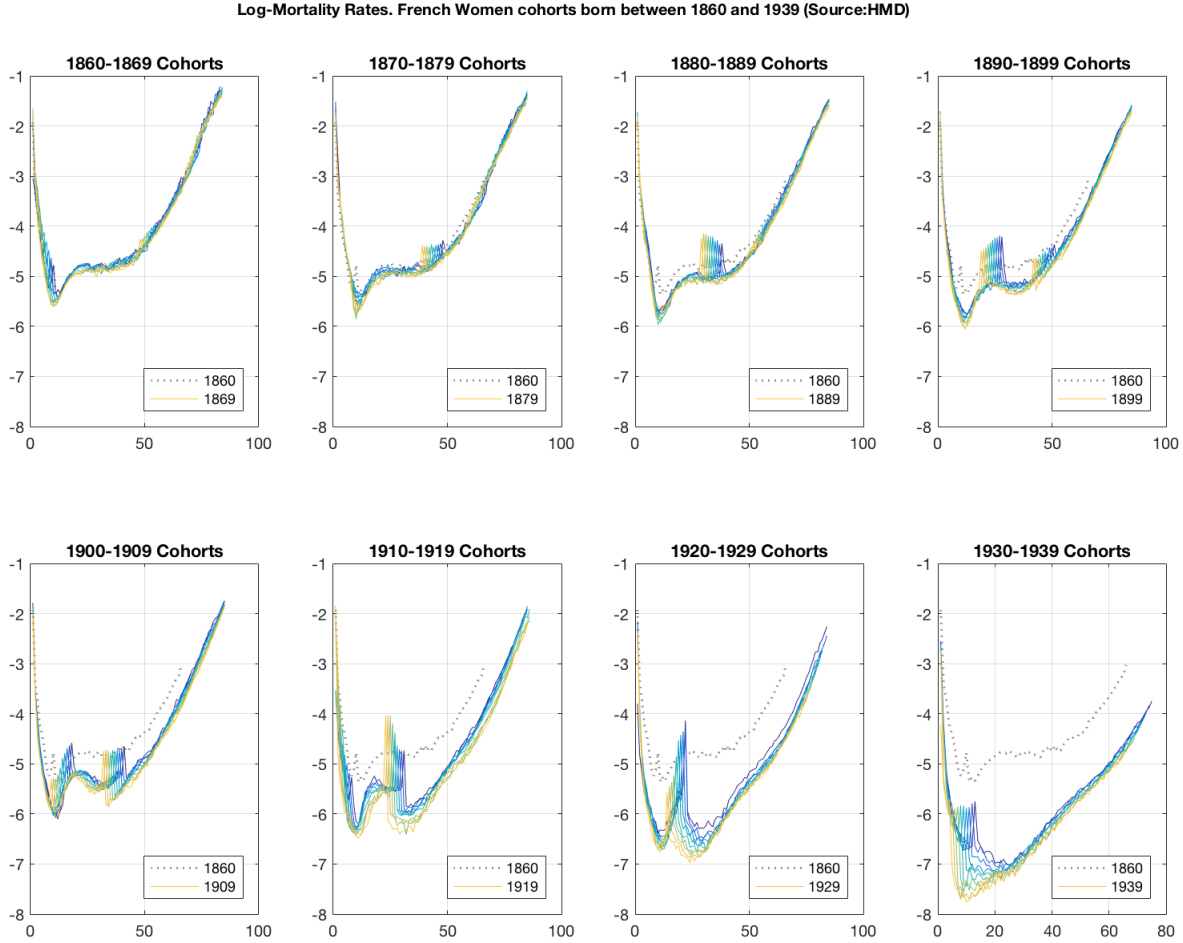

Figure 20: Age profile of mortality of men born in France between 1860 and 1940, by decade

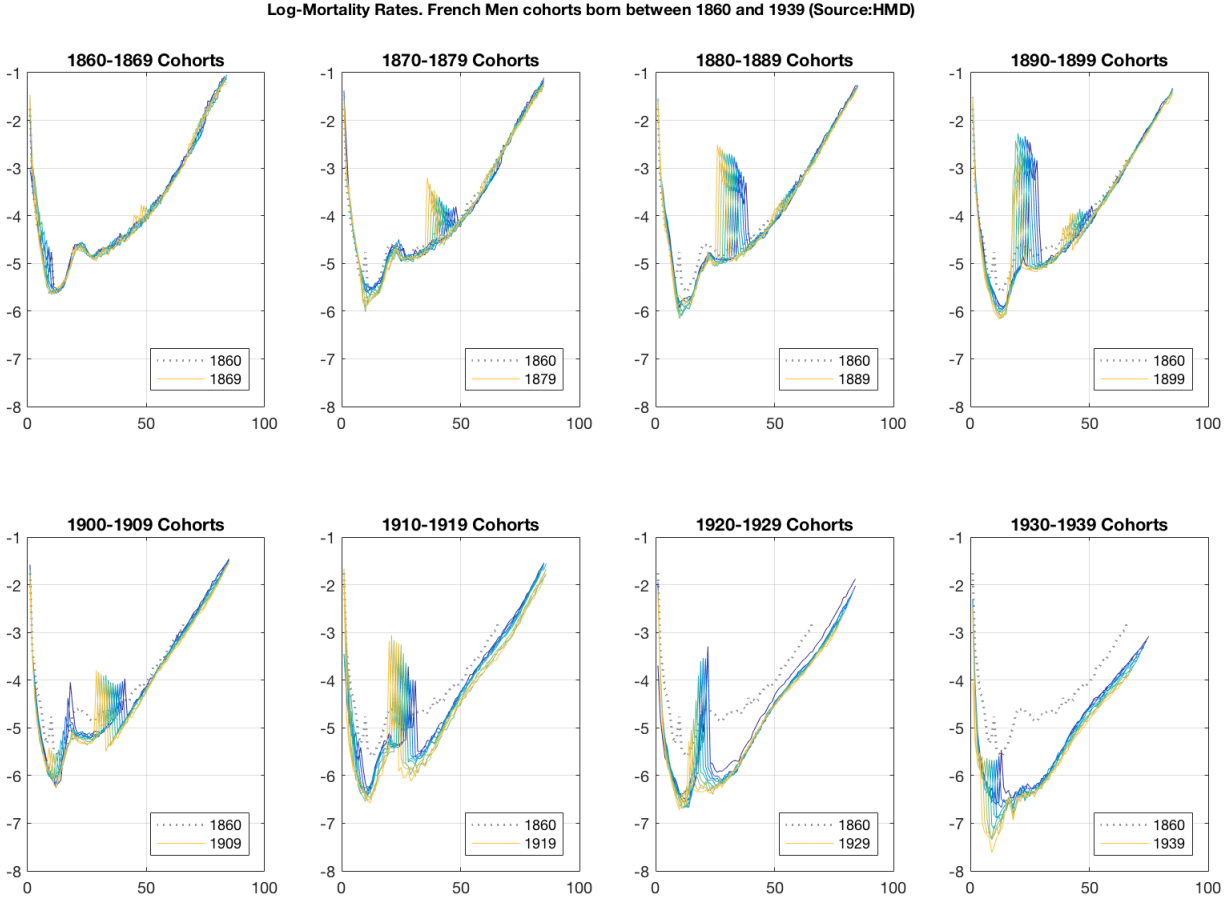

Figure 21: Health and mortality in the first two years of life

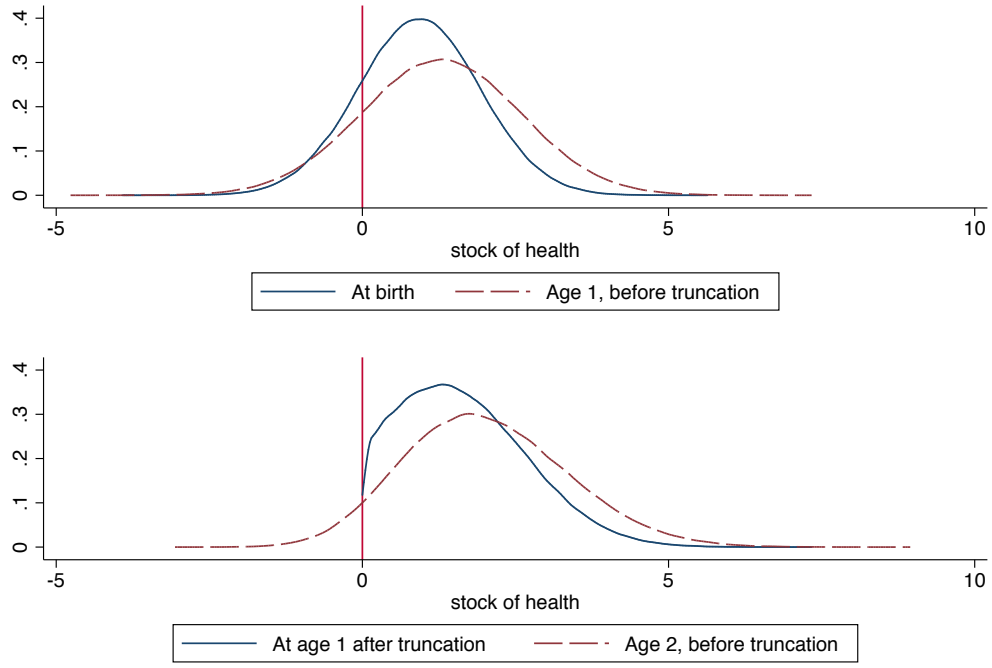

Data from simulations

In the first period the (infant) mortality rate MR_1 is given by

$$MR_1 = P(H_1 \leq \underline{H}) = P\left(H_0 + I - \delta + \varepsilon_1 \leq \underline{H}\right)$$

$$= P(\varepsilon_1 \leq \varphi_1) = F(\varphi_1)$$

where $\varphi_1 = \underline{H} - I + \delta - H_0$ captures the threshold for dying in period 1 in terms of the random shock. Investments lower this threshold (lower mortality) and depreciation increases it (increases mortality).
Consider now the probability of dying at age $t = 2$. This is given by the probability that the stock falls below $\underline{H}$ at age 2, conditional on having survived to age 2, which can be expressed as:

$$\begin{aligned} MR_2 &= E(D_2 = 1 | D_1 = 0) = P(H_2 < \underline{H} | H_1 > \underline{H}) \\ &= \frac{P(H_2 < \underline{H}, H_1 > \underline{H})}{P(H_1 > \underline{H} | g_1, g_2)} = \frac{P(\varepsilon_2 < \varphi_2 - \varepsilon_1, \varepsilon_1 > \varphi_1)}{1 - F(\varphi_1)} \\ &= \frac{K(\varphi_2, \varphi_1)}{1 - F(\varphi_1)} \end{aligned} \tag{18}$$

where $\varphi_2 = \underline{H} - I + \delta 2^{\alpha} - H_0$ captures the threshold for dying in period 2, and $K(\varphi_2, \varphi_1) = \int_{\varepsilon_1=\varphi_1}^{\infty} \int_{\varepsilon_2=-\infty}^{\varphi_2-\varepsilon_1} f(\varepsilon_1) f(\varepsilon_2) d\varepsilon_1 d\varepsilon_2$ is the density right above the old threshold and below the new threshold, that is the fraction of survivors who dies as a result of a new shock. The denominator is the fraction of survivors.

Figure 22: Comparative statics for health

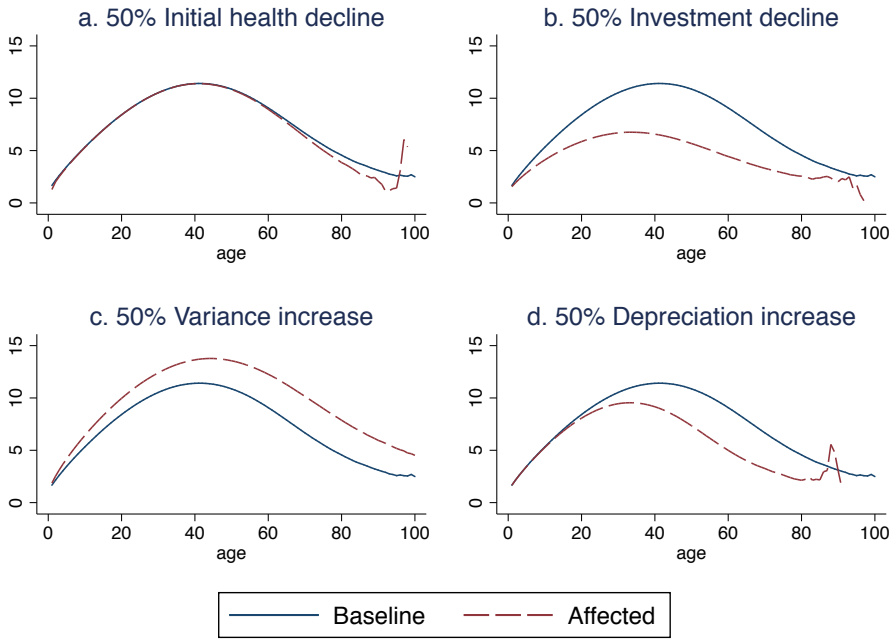

Note: Simulated data for a population of 500,000 individuals. The figures show the effect of changes relative to the baseline model, which is simulated using the same parameters we used for Figure 2.

Figure 23: Effects of variance increase

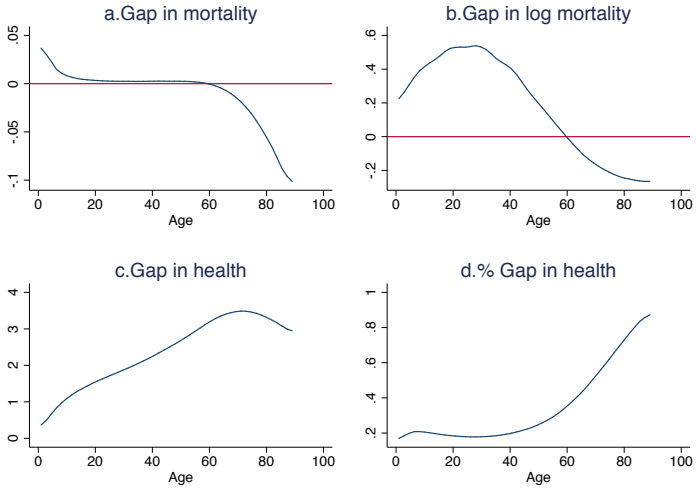

Note: Simulated data for two population of 500,000 individuals each. The figures show the effect of changes relative to the baseline model, which is simulated using the same parameters we used for Figure 3.

Figure 24: Effect WWI and WWII on female mortality

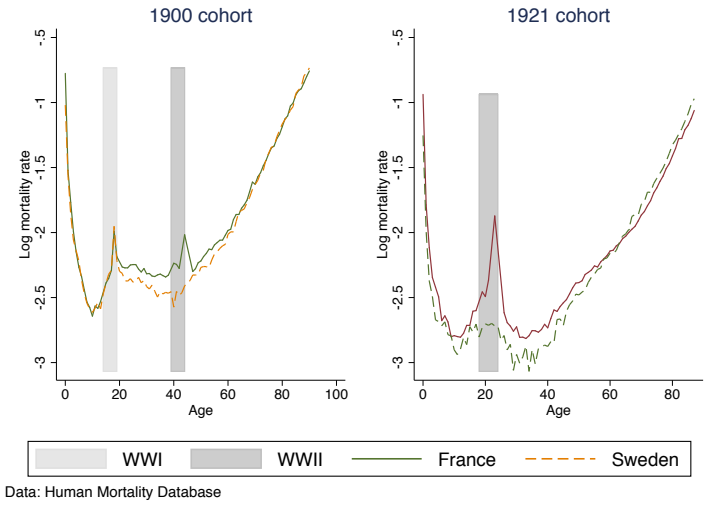

Data: Human Mortality Database

Figure 25: Female population in France since 1816

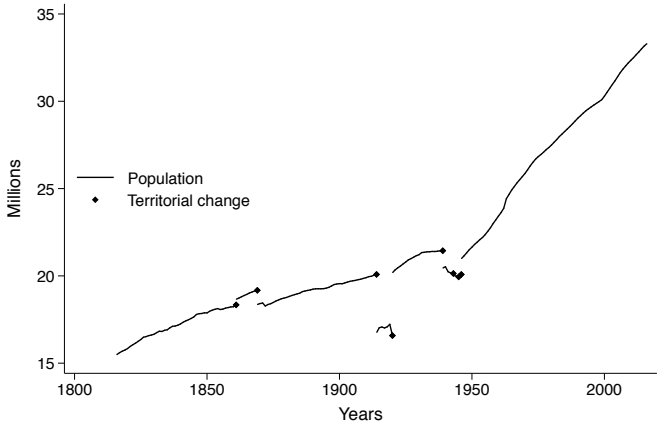

Note: See the technical documentation of the Human Mortality Database for details about the population coverage for the French mortality data.

Figure 26: Comparison of q-rate in the paper and HMD (1816)

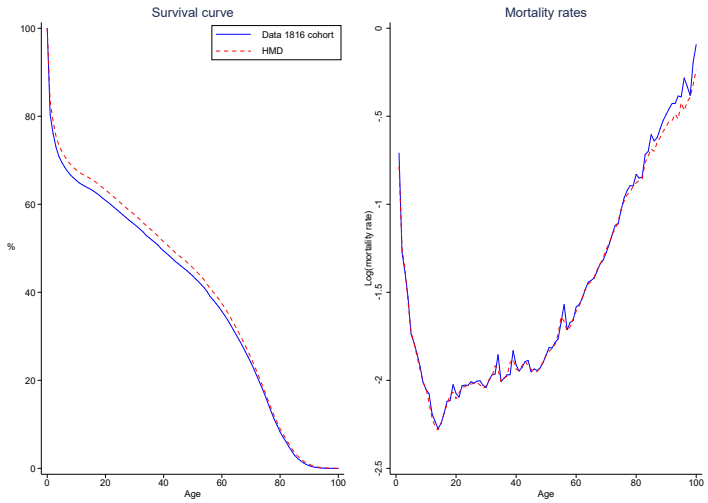

Life expectancy: 38.25 (with the q we use) and 39.86 (with the q in HMD). The life expectancy in HMD is 39.83.