The Lifetime Costs of Bad Health

Mariacristina De Nardi, Svetlana Pashchenko, and Ponpoje Porapakkarm*

March 18, 2022

Abstract

What generates the observed differences in economic outcomes by health? How costly it is to be unhealthy? We show that health dynamics are largely driven by ex-ante fixed heterogeneity, or health types, even when controlling for one's past health history. In fact, health types are the key driver of long spells of bad health. We incorporate these rich health dynamics in an estimated structural model and show that health types and their correlation with other fixed characteristics are important to account for the observed gap in economic outcomes by health. Monetary and welfare losses due to bad health over the life cycle are large, concentrated, and to a large extent due to factors pre-determined earlier in life. A large portion of the related monetary costs is due to income losses, especially for people of working age, while a substantial portion of the welfare losses arises because health affects life expectancy.

JEL Codes: D52, D91, E21, H53, I13, I18

Keywords: health, health insurance, medical spending, wealth-health gradient, life-cycle models

*De Nardi: University of Minnesota, Federal Reserve Bank of Minneapolis, CEPR, and NBER. Pashchenko: University of Georgia. Porapakkarm: National Graduate Institute for Policy Studies (GRIPS). We thank Liran Einav, Christopher Flinn, Eric French, Joseph Altonji, Selahattin Imrohoroglu, John Kennan, Naoki Aizawa, Sagiri Kitao, Rasmus Lentz, Vincenzo Quadrini, Yongseok Shin, Matthew White, and all seminar participants at the Keio University, University of Hong Kong, Chinese University of Hong Kong, University of Tokyo, University of Delaware, University of Connecticut, FED Chicago, FED Atlanta, Institute for Fiscal Studies, AMES2017, SED2017, EMES2017, SAET2018 for their comments and suggestions. De Nardi gratefully acknowledges support from the ERC, grant 614328 "Savings and Risks." Porapakkarm gratefully acknowledges support from the JSPS KAKENHI Grant Number 15K03505. Pashchenko acknowledges financial support through the Terry-Sanford Research Award. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research, the CEPR, any agency of the federal government, or the Federal Reserve Bank of Chicago.

1 Introduction

There are substantial differences in economic outcomes by health. For example, even among the relatively homogeneous group of men with a high-school degree, the healthy earn, on average, 37% more than the unhealthy conditional on working. The difference in wealth is even more remarkable. The gap in wealth by health starts at a relatively young age and becomes very large by retirement time: in the same group, the median wealth of the healthy at age 65 is 65% larger than that of the unhealthy.[1]

These facts raise two important questions. First, what generates such a large difference in economic outcomes by health? Second, given these large and prolonged differences, how costly it is to be unhealthy from the entire life-cycle perspective?

We address these questions using a structural framework. In general, the differences in economic outcomes by health can be due to three mechanisms: (i) health changes economic circumstances; ii) economic circumstances affect health, (iii) healthy and unhealthy people are innately different. While we adopt the first mechanism, the novelty of our approach is that we also allow for the third mechanism. In fact, modeling the latter is where we see a big gap in the literature. While the existence of the third mechanism is well-recognized, most structural models adopt a simple structure in which all innate differences across people are entirely captured by an observable and unidimensional variable, such as education.

Hence, a related fundamental, and still unanswered question is what is the role of fixed unobserved characteristics—for instance, characteristics driven by genetics and experiences early in life—and what are their long-lasting effects on health and economic outcomes. A convincing answer to this question necessitates a much richer modeling of innate differences, which are likely multidimensional, and possibly correlated with each other.

Our focus on the role of factors that are predetermined early in life, or at birth, goes in parallel with the recently growing empirical literature that documents the importance of the effects of early-in-life factors over the entire life course (Anda et al., 2006; Barth et al., 2020; Case et al. 2005; Conti et al., 2005; Cronqvist and Seigel, 2015; Felitti et al., 1998; Harris et al., 2017). Because the findings from this literature are very compelling, it is crucial to incorporate these new forces in structural models. Our paper does so in the context of a rich quantitative framework.

Our analysis proceeds in several steps. We document several novel facts about long-run health dynamics showing that they are complex and not consistent with a low-order Markov process. Specifically, health transitions display strong duration dependence: the longer people have been unhealthy (healthy), the less likely they are to become healthy

[1] Own calculations, see Section 3.4 for details.

(unhealthy). We formulate and estimate a parametric model of health shocks that allows for both history-dependence and fixed ex-ante heterogeneity, and that matches the patterns that we observe. Finally, we incorporate our estimated health process in a life-cycle model. The innate differences across individuals in our model are three-dimensional and include fixed heterogeneity in health (or health types), in labor productivity, and in the rate of time preferences. We estimate the first two fixed factors and their correlation together with our health process. We estimate the correlation between health types and the rate of time preferences in our structural model by using the Method of Simulated Moments.

We focus on high school men to avoid the confounding effect of education and gender on health and economic outcomes and to emphasize the role of fixed heterogeneity within this relatively homogeneous group. For more accurate measurement, we use three datasets: the Health and Retirement Study, the Panel Study of Income Dynamics, and the Medical Expenditure Panel Survey.

Our estimated model is consistent with three sets of important facts. First, it captures the dynamics of health, including its duration dependence. Second, it matches the observed impact of bad health on earnings and labor supply (income-health gradient), medical spending, and life expectancy. Third, it captures the wealth-health gradient by matching the large difference in wealth levels between the healthy and unhealthy over the lifespan.

Our first set of results relates to the estimation of the health process. We find that both history-dependence and fixed ex-ante heterogeneity (or health types) are important drivers of health dynamics and that the role of health types is significant even when we control for long history dependence. In addition, and importantly, the variation in the transition probabilities across health types is much larger than that across health histories. This implies that health types play an important role in explaining the high persistence of health and the occurrence of long spells of bad health. Specifically, health types account for more than 70% of the variation of the fraction of lifetime spent being unhealthy. We also show that fixed heterogeneity in one's health shock process is correlated with fixed heterogeneity in labor productivity: people with low fixed productivity are more likely to be of a worse health type.

Our second set of results relates to the implications of our estimated structural life-cycle model. Our estimates imply a strong correlation between one's health type and rate of time preferences: among the long-term unhealthy a larger fraction is less patient and has a lower propensity to save. This is important to account for the wealth-health gradient: when the correlation between patience and health type is shut down, the model substantially underpredicts the wealth gap between the healthy and unhealthy, even though it still matches the income-health gradient.

We find that the lifetime costs of bad health are large when measured both in monetary

and welfare terms. On average, people lose about $1,500 per year over their entire life because of bad health. Our measure of monetary costs includes both direct (out-of-pocket medical spending) and indirect (loss in labor earnings) costs. We find that the latter component is a large contributor to the lifetime costs of bad health, especially for working-age people, and arises because unhealthy individuals are less productive and work less than healthy ones. In fact, even though total medical costs are substantial for the long-term unhealthy, the effects of out-of-pocket costs are smaller due to health insurance coverage.

Our welfare measure of the lifetime costs of bad health represents the consumption equivalent variation for individuals in the counterfactual situation when they (unexpectedly) never draw bad health realizations. We find that moving from this counterfactual to the baseline would be equivalent to losing 10% of annual consumption, on average. A decomposition exercise shows that the major contributor to the welfare costs is the effect of bad health on expected lifespan: because life is valuable, bad health reduces welfare by shortening the length of life.

Finally, we show that the monetary and welfare costs of bad health are very concentrated and highly unequally distributed across health types. This happens because people with different health types are strikingly different in terms of the fraction of their lifetime that they spend being unhealthy. For example, on average, people with the worst health type spend almost two third of their lifetime in bad health. This translates into a substantial contribution of health types to the variation in the lifetime costs of bad health. That is, health types account for close to 70% for the monetary costs and about a third for welfare costs due to bad health. The smaller contribution of health types to the variation in welfare costs is due to the fact that welfare costs are very sensitive to lifespan length and that health types affect the length of unhealthy spells much more than they affect life expectancy.

Our study thus provides several novel contributions. First, it documents new facts related to the long-run dynamics of health and shows that to account for this, health process has to feature both long memory and fixed health types.

Second, it shows that the correlated multidimensional ex-ante heterogeneity in health, labor productivity, and patience to a large extent contributes to the joint evolution of health, income, and wealth over the life cycle. Thus, to account for the observed disparities in economic outcomes by health in a life-cycle model, it is important to take ex-ante heterogeneity seriously. We thus add to the existing studies that commonly attribute the observed health-related disparities in economic outcomes to ex-post health shocks and economic circumstances during adulthood.

Third, we develop a quantitative model designed to gauge the long-run or accumulated effects of bad health over the entire life cycle. Using our framework, we are able to measure

the comprehensive effects of bad health both in terms of monetary and welfare costs, and evaluate the role of factors pre-determined early in life in generating these costs. To the best of our knowledge, the accumulated effects of bad health over the life cycle have not been assessed before due to the lack of data or an appropriate structural framework.

Related literature. The robust relationship between health and economic outcomes is well-documented in the literature (see Cutler et al., 2011 for a review). There is growing interest in using structural models including health to study a broad set of issues such as health insurance reform (see Fang and Krueger, 2021, for an extensive review) or the contribution of heath to inequality in earnings (Capatina et al., 2021; Hosseini et al., 2021b). The findings of these studies crucially depend on what generates difference in economic outcomes by health: a direct link between the two or innate differences across individuals. We contribute to the literature by emphasizing the role of correlated fixed characteristics for the systemic economic disparity between health groups.

At a methodological level, quantitative structural studies on health-related questions emphasize the importance of three points: (i) how we measure health, (iii) how best to describe health dynamics, (iii) how to model the direct link between health and economic outcomes.

In our study, we use self-reported health status, a discrete measure of health. This variable is most commonly used in structural work partly because it is available in several micro datasets and is consistently measured across them.[2] In addition, a number of studies find that self-reported health is highly correlated with other subjective and objective measures of health and also has significant explanatory power in predicting future mortality, even after controlling for many other factors (see Idler and Benyamini (1997) for a review, and van Doorsaler and Gerdtham (2002), and Pijoan-Mas and Ríos-Rull (2014) for a more recent examination).

Several recent studies have suggested an alternative measure of health: a continuous index constructed by aggregating several variables such as diagnosis of most common health conditions, limitations of activities of daily living, cognitive impairments, etc. (Blundell et al., 2020a; Hosseini et al., 2021a; Poterba et al., 2017). Capatina et al. (2021) and Prados (2018) construct a more complex measure of health based on an extended list of detailed diagnosed medical conditions available in the Medical Expenditure Survey (MEPS) Dataset. Such a measure is arguably closer to the objective underlying health, but the short panel dimension of the MEPS offers limited opportunities to study its long-term dynamics.

[2] We show that this variable is consistently measured in the PSID and the HRS (see Figures 1-2 in the next section). Attanasio et al. (2011) compare this variable in the HRS and MEPS, and shows that the two datasets are consistent.

Importantly, Blundell et al. (2020b) point out that properly accounting for innate differences across individuals is a more important modeling issue than the choice of health measure. Our estimation provides new insights about the role of these innate differences.

Turning to the dynamics of health, structural studies commonly assume it follows a first-order Markov process. Some studies using a continuous health measure assume the AR(1) specification augmented with transitory shocks and/or fixed effects (Blundell et al., 2020a; Hosseini et al. 2021b). One of the key contribution of our study is to show that health process has a long memory.

Regarding the mechanism *directly* linking health and economic outcomes, structural studies follow one of the following approaches. The most common approach is to assume that health is exogenous and affects medical spending, labor productivity, and other economic variables (Conesa et al. 2018; Kitao and Jeske, 2009; Pashchenko and Porapakkarm, 2013, 2017). The second approach is to assume that health is endogenous and can be affected by medical spending (Jung and Tran, 2016; Ozkan, 2018), effort (Cole et al., 2018) or healthy/unhealthy habits (Bolt, 2021; Hai and Heckman, 2021).[3] In modeling the direct link between health and economic outcomes, we follow the first approach, but we emphasize the *indirect* link coming from the fact that both health and economic variables are affected by correlated fixed factors.

At an applied level, there is a long list of empirical studies documenting the impact of health on labor marker outcomes (see Currie and Madrian, 1999, for a review). A number of structural studies investigate the importance of health for economic decisions such as savings (Ameriks et al., 2020; De Nardi et al., 2010, 2016; Hubbard et al., 1995; Lockwood, 2018; Nakajima and Telyukova, 2020) or labor supply (Capatina, 2015; French, 2005; French and Jones, 2011). However, the total losses imposed by bad health over the entire life-cycle has not been assessed before. Arguably, fully understanding the accumulated costs of bad health, its sources, and to what extent these costs are predetermined by factors formed at birth or early age is very important for analyzing any health-related policy issues.

It is also important to mention the relationship to the literature assessing the value of changes in longevity from the perspective of an individuals' maximization of expected lifetime utility. Early examples include Arthur (1981), Rosen (1988), Shepard and Zeckhauser (1984). More recently, Murphy and Topel (2006) apply this approach to quantify the welfare gains from increased longevity during the 20th century and find that they are substantial. These studies typically abstract from individual-level heterogeneity and focus on how

[3] A somewhat mixed approach is used by De Nardi et al. (2016) and Pashchenko and Porapakkarm (2019) who assume medical spending are (partially) endogenous but cannot affect health. Another hybrid approach is used by Capatina et al. (2021) who assume (endogenous) employment status can affect heath transitions.

a representative cohort values longevity increases. We contribute to this line of research by evaluating the welfare losses from shorter life expectancy arising from health shocks of individuals who differ in ex-ante fixed characteristics and by decomposing these costs.

The rest of the paper is organized as follows. Section 2 documents empirical facts related to health dynamics and estimates the health shock process. Section 3 introduces our life-cycle model and Section 4 describes its estimation. We present the results and conclusions in Section 5 and Section 6, respectively.

2 Health and labor productivity

For our health process estimation, we primarily use the Panel Study of Income Dynamics (PSID). When possible, we also use the Health and Retirement Study (HRS) to cross-check moments from the PSID and to validate our estimated model. We report more detail about these data sets, our samples, and how we use them in Appendix A.

For each data set, we use a sample of male household heads with 12 to 14 years of education (corresponding to the high school degree or at most 2 years of college education). We normalize all nominal variables to the base year (2013) using the Consumer Price Index (CPI).

We start by documenting the cross-sectional and time-series properties of self-reported health status and individual-level fixed labor productivity. We then estimate a rich process for health dynamics that is consistent with these moments and discuss its implications.

2.1 The evolution of health

In the PSID and the HRS, people rank their health as *excellent*, *very good*, *good*, *fair* or *poor*. As common in the literature, we aggregate these answers into a binary measure of health (see, for example, French (2005), Capatina (2015), Rust and Phelan (1997)) and classify as *healthy* or in *good health* people who report their health to be in the first three categories, and as *unhealthy* or in *bad health* people who report being in fair or poor health.

The left graph of Figure (1) displays the percentage of people that we classify as in bad health by age. The dots correspond to the PSID data, while the crosses refer to HRS data. The center and right panels display the two-year health transition probabilities. Details of their construction are in Appendix B.1.

Consistent with previous findings, more people are in bad health when older, the probability of recovering from bad health decreases with age, and that of becoming sick increases with age. In addition the PSID and the HRS yield a very similar picture.

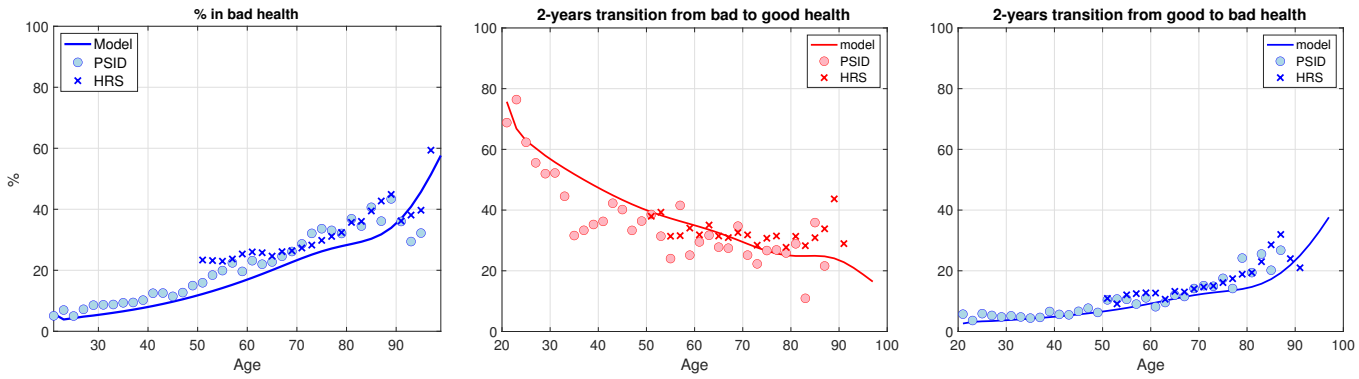
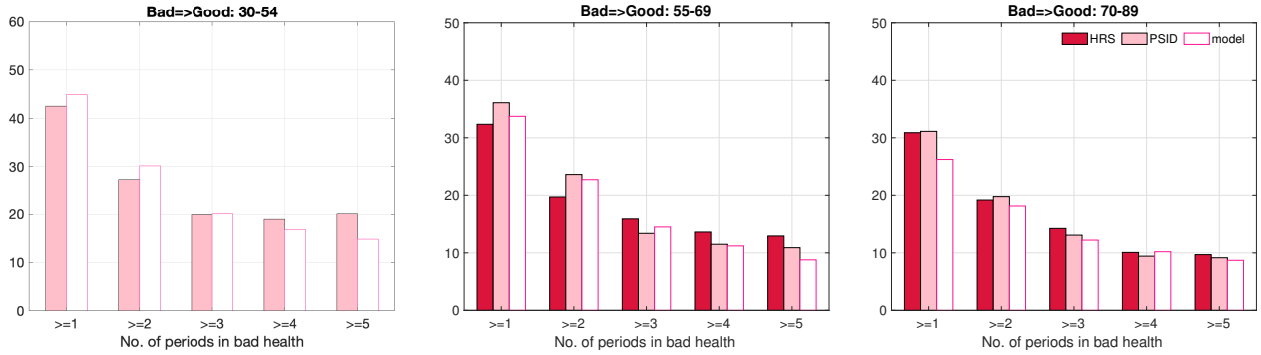


Figure 1: Health. Left panel: percentage of people in bad health by age. Middle panel: percentage of people moving from bad to good health in the next two years. Right panel: percentage of people moving from good to bad health in the next two years.

a. Percentage of transitions from bad to good health conditional on being in bad health



b. Percentage of transitions from good to bad health conditional on being in good health

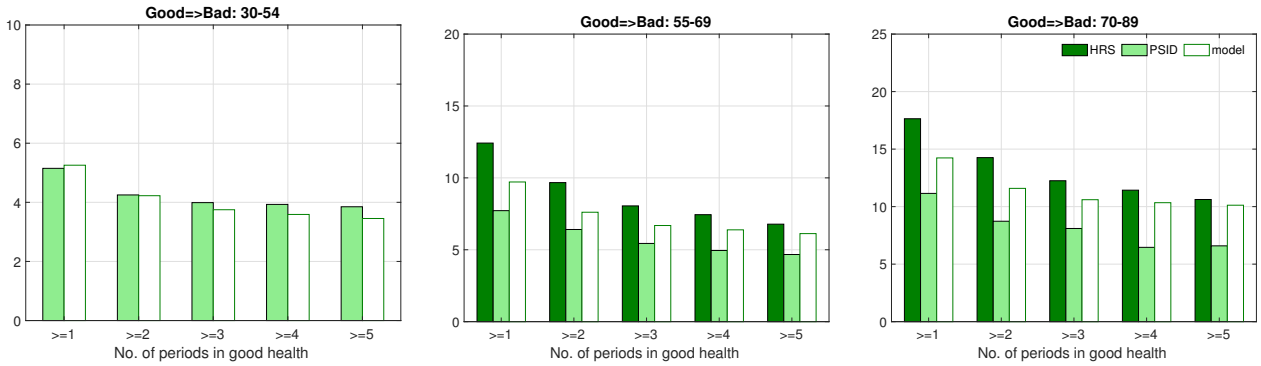


Figure 2: Dynamics of health conditional on duration of previous health status. Each period is every two years.

To better understand the dynamics of health, the top and bottom panel of Figure 2 display the probabilities of moving from bad (good) to good (bad) health over the next two years, conditional on being in bad (good) health for at least τ consecutive periods (with a period being two years). We group observations in three age groups: 30-54, 55-69, and older

than 70 and report the resulting statistics from left to right.

An important feature of these graphs is that both the probability of recovering from bad health and that of becoming unhealthy decline monotonically with duration: the longer an individual has been unhealthy (healthy), the less likely he is to become healthy (unhealthy). This pattern holds for all age groups.[4]

This form of duration dependence cannot be captured by the low-order Markov processes for health that are commonly used in the literature (e.g., French, 2005; French and Jones, 2011; and Capatina, 2015). In fact, a first order Markov process implies that the transition probability does not depend on how long one has been in the current health state. A second order Markov process implies that one's transition probability is the same for durations longer or equal to two periods.

This negative duration dependence can be generated by two mechanisms. First, the effects of health can compound. People who are sick for a long period of time might have a lower recovery probability than those who are sick for a shorter period of time. Similarly, people who recently recovered from bad health might have a larger probability of relapsing than those who have been in good health for longer. This mechanism can be well represented by a high-order Markov process.

Second, people may differ in their ability to recover or in their predisposition to become sick. In this case, people who are more likely to recover move out of the bad health state faster. Hence, the pool of the long-term unhealthy is predominantly composed of individuals who are inherently less likely to recover. Likewise, among the pool of healthy people, some might be more susceptible to a certain health condition, for instance due to their genetics or lifestyle. This second mechanism is consistent with fixed heterogeneity in health transition probabilities. In formulating our health process, we allow for both mechanisms (Section 2.3).

2.2 Health status and labor productivity

Previous work documents large heterogeneity in unobserved fixed characteristics that affect people's labor market outcomes, even within the same education group. Among others, see Guvenen (2009) and Capatina (2015) for estimations that adopt different specifications of labor income processes. Our goal now turns to determining whether fixed labor productivity is related to health and its dynamics.

To compute fixed labor productivity, we start by computing real labor income from the PSID for workers younger than 70, who work at least 520 hours per year, and earn at least the

[4] This negative duration dependence is a robust pattern. It also holds when we use the annual data (using the 1984-1997 waves of the PSID). See De Nardi et al. (2018) for those results. It is also present when we exclude people ever receiving Social Security Disability Insurance.

federal minimum wage. As French (2005), we then estimate the age profile of labor income by health status among working individuals using the following fixed-effect regression:

$$log(inc_{it}) = \sum_{t=21}^{69} \sum_{j=\{G,B\}} d_t^j \times D_{it}^{age} \times D_{h_{it}=j} + \widehat{\gamma}_i + u_{it}, \tag{1}$$

where inc_{it} is persons' i labor income at age t, d_t^j it the coefficients a dummy variable capturing the interaction between one's age and health status, and $\widehat{\gamma}_i$ is unobserved fixed labor productivity. Since this specification includes cohort effects in individual fixed productivity, we subsequently regress individual-level estimated fixed productivity on birth-year dummies to remove them. We denote the resulting estimates as γ_i.

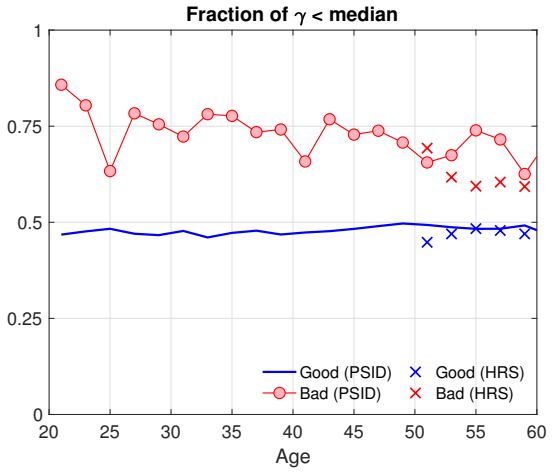


Figure 3: Fixed wage productivity by health status for both workers and non-workers: fractions of individuals with fixed productivity below the median (solid line for the healthy and line with dots for the unhealthy).

Figure 3 reports the fraction of people whose fixed labor productivity (γ_i) is below the median for the healthy and the unhealthy. It shows that, for all age groups, the fraction of unhealthy people having γ_i below the median (dotted line) is much higher than 50%, while the corresponding fraction among the healthy (solid line) is about 50%. The results from the HRS (crossed line) are very similar. This indicates that the composition of fixed labor productivity of the healthy and unhealthy is different and that we need to account for this.

2.3 Our health process

From the previous sections, we have learned that health dynamics are not well represented by low-order Markov processes and that a larger fraction of the unhealthy have low fixed

labor productivity. In this section, we formulate a process for the evolution of health that can capture both observations.

To account for the first observation, we allow our health shock process to feature both fixed heterogeneity and a long history dependence. One important consideration, however, is that the negative duration dependence that we find could partly result from the fact that we combine poor and fair health in our measure of bad health, but that, in reality, these two health states may differ in their persistence. To address this possibility, we no longer merge poor and fair health into "bad" health. That is, we allow for three health states in our estimated health process: $\{P, F, G\}$ for *poor, fair*, and *good* health, respectively.

To account for the second observation, we allow one's fixed health type to be correlated with one's fixed ex-ante heterogeneity in labor productivity.

Furthermore, we set up our process for health dynamics based on two criteria. First, it must capture the cross-sectional and dynamic moments of health that we document. Second, it must be parsimonious, so that a structural life-cycle model including this health process is computationally manageable.

As a result, we formulate our 3-state health shock process as an ordered-logit model with fixed unobserved heterogeneity. We also assume that the future evolution of health depends on an age polynomial, whose coefficients depend on which of the three health states the person is during the current period. In terms of longer memory dependence, we restrict the evolution of health to depend on the number of consecutive periods of previous good or bad health, that is, we combine the states of poor and fair health. In addition of preserving tractability, combining the history of poor and fair health avoids the dealing with the problem that the number of people being in poor health for several consecutive waves is small.

More specifically, we assume that for an individual with current poor or fair health ($h_t \in \{P, F\}$), who has been in bad health for τ_B periods, the probability of being in poor health, conditional on surviving to age $t+1$, follows a logistic function (for the ease of notation, we suppress individuals' level subscript i)

$$Pr\left(P_{t+1} \mid h_t, \tau_B, \eta\right) = \mathbf{\Lambda}\left(\sum_{\tau=1}^{T-1} a_\tau^B \mathbf{1}_{(\tau_B=\tau)} + a_T^B \mathbf{1}_{(\tau_B \geq T)} + f_{age}^{h_t}(t) + a_\eta^B D_\eta\right), \tag{2}$$

and that the probability of being in either fair or worse health at age $t+1$ is

$$Pr\left(F_{t+1} \cup P_{t+1} \mid h_t, \tau_B, \eta\right) = \mathbf{\Lambda}\left(\sum_{\tau=1}^{T-1} a_\tau^B \mathbf{1}_{(\tau_B=\tau)} + a_T^B \mathbf{1}_{(\tau_B \geq T)} + f_{age}^{h_t}(t) + b_1 + a_\eta^B D_\eta\right), \tag{3}$$

where $\mathbf{\Lambda}$ is the logistic function, and $\mathbf{1}_{(\cdot)}$ is an indicator function which is equal to one if its

argument is true and zero otherwise. The first two terms in the bracket capture the history dependence, where τ_B is the number of consecutive periods an individual has been in bad health (either poor or fair). We denote as T the longest possible history dependence (i.e., how many lags, including the current period, are included in our process).

We allow the transition probability to depend on age through $f^P_{age}(t)$ and $f^F_{age}(t)$, which are second-order polynomials in age. Their coefficients can differ depending on whether one's current health status is poor or fair. The term b_1 is a non-negative constant to ensure that the cumulative probability is monotonically increasing.

The last term in the logit function, D_η, is the dummy variable for one's fixed *health type*. We assume there are three such types, which we denote as $\{\eta_1, \eta_2, \eta_3\}$, and impose the following rank in our estimation: $a^B_{\eta_1} > a^B_{\eta_2} > a^B_{\eta_3}$. Thus, η_1-type has a lower chance to be in better health than η_2- and η_3-types. We will refer to an individual with η_1 and η_3 as the worst and best health types, respectively.

Similarly, the probability of being in poor health (and probability to be in poor or fair health), conditional on surviving to age $t+1$ and being healthy for τ_G periods, is

$$Pr\left(P_{t+1} \mid G_t, \tau_G, \eta\right) = \mathbf{\Lambda}\Big(\sum_{\tau=1}^{T-1} a^G_\tau \mathbf{1}_{(\tau_G=\tau)} + a^G_T \mathbf{1}_{(\tau_G \geq T)} + f^G_{age}(t) + a^G_\eta D_\eta\Big), \tag{4}$$

$$Pr\left(F_{t+1} \cup P_{t+1} \mid G_t, \tau_G, \eta\right) = \mathbf{\Lambda}\Big(\sum_{\tau=1}^{T-1} a^G_\tau \mathbf{1}_{(\tau_G=\tau)} + a^G_T \mathbf{1}_{(\tau_G \geq T)} + f^G_{age}(t) + b_2 + a^G_\eta D_\eta\Big), \tag{5}$$

where the type-dependent coefficients a^G_η allow each health type to have a different effect on the transition probabilities of people currently in good and bad health.

It is worth noting that our specification nests the first-order Markov process for health shocks that is commonly used in the literature, when we restrict: $T = 1$ and $a^B_\eta = a^G_\eta = 0$ for all health types.

Since one's health type η is unobservable, we model the probability of an individual having a certain health type as an ordered-logit model. Hence, the cumulative probability of an individual having health type $\eta \leq \eta_j$ can be expressed as follows:

$$Pr(\eta \leq \eta_j \mid \mathbf{X}^\eta_{t_0}) = \begin{cases} \mathbf{\Lambda}\left(b_{\eta_j} - \mathbf{B}_\eta \times \mathbf{X}^\eta_{t_0}\right) & \text{for } \; j = \{1,2\}, \\ 1 & \text{for } \; j = 3, \end{cases} \tag{6}$$

where b_{η_1} and b_{η_2} are the constant terms in the ordered-logit model, with $b_{\eta_1} < b_{\eta_2}$, t_0 is the earliest age at which we observe an individual in our sample, $\mathbf{X}^\eta_{t_0}$ is a set of characteristics that can be informative about one's health type and are observable as of at age t_0, and $\mathbf{B}_\eta$ are the corresponding coefficients (we provide more details on the construction of the

log-likelihood function in Appendix B.2).

The variables that we include in $\mathbf{X}_{t_0}^{\eta}$ are age t_0, birth-year dummies (in 10-year brackets), health status h_{t_0}, fixed labor productivity γ, and net worth k_{t_0}. We allow for tercile dummies for fixed productivity $(\gamma_L, \gamma_M, \gamma_H)$, and hence for a non-linear relationship between γ and η. We allow for quintile dummies for net worth, for each 5-year age bracket. The next section reports our estimation results and discusses why we include these variables and what we learn from our estimated processes.

Next, we turn to modeling survival. We do so by specifying the following logit model for the two-year survival probability of an individual at age t, who has been in health state h_t for τ_h periods,

$$\begin{aligned}
\zeta_t\left(h_t, \tau_B\right) =& \mathbf{\Lambda}\left(\sum_{\tau=1}^{2} a_{\tau}^{\zeta B} \mathbf{1}_{(\tau_B=\tau)} + a_3^{\zeta B} \mathbf{1}_{(\tau_B \geq 3)} + f_{age}^{\zeta h_t}(t)\right) \qquad \text{if} \quad h_t \in \{P, F\}, \\
\zeta_t\left(h_t, \tau_G\right) =& \mathbf{\Lambda}\Big(\sum_{\tau=1}^{2} a_{\tau}^{\zeta G} \mathbf{1}_{(\tau_G=\tau)} + a_3^{\zeta G} \mathbf{1}_{(\tau_G \geq 3)} + f_{age}^{\zeta G}(t)\Big) \qquad \text{if} \quad h_t = G. \qquad (7)
\end{aligned}$$

The first two terms in the logit function capture history dependence. Similar to the health shock process, we combine the history of poor and fair health into that of bad health. The terms $f_{age}^{\zeta h_t}(t)$, where $h_t \in \{P, F\}$, and $f_{age}^{\zeta G}(t)$ are linear functions of age.

Since the PSID's sample size for older people is small, we estimate our survival probabilities using data on males with a high school degree from the HRS. We then extrapolate them to obtain the survival probabilities for younger people. Figure B1 in Appendix B.3 shows our estimated two-year survival probabilities.

After estimating our survival probabilities, we estimate Eqs (2)-(6) jointly by maximum likelihood. We report the results in the next section. In estimation, we do not impose any restrictions other than those necessary for a standard ordered-logit model. That is, in Eqs (2)-(3) we impose $a_{\eta_1}^B > a_{\eta_2}^B > a_{\eta_3}^B$ where $a_{\eta_2}^B$ is normalized to zero. Similarly, in Eqs (4)-(5) $a_{\eta_2}^G$ is normalized to zero but there are no restrictions on $a_{\eta_1}^G$ and $a_{\eta_3}^G$. Thus, we allow for any possible ranking among $\left\{a_{\eta_1}^G, a_{\eta_2}^G, a_{\eta_3}^G\right\}$. Finally, a cumulative probability must be monotonically increasing; thus, b_1 in Eq. (3) and b_2 in Eq. (5) must be non-negative, and $b_{\eta_1} \leq b_{\eta_2}$ in Eq. (6) .

2.4 Estimation results

We estimate several versions of our health process by allowing for history dependence T from 1 to 5 lags. We report the results in Tables 1 and 2. Two results are worth noting. First, in all specifications, the coefficients on health types $\left(a_{\eta_1}^B, a_{\eta_3}^B, a_{\eta_1}^G, a_{\eta_3}^G\right)$ are statistically

significant. Hence, even when we allow for 5 periods (corresponding to 8 lagged years) of history dependence, health types matter for health evolution. This suggests that our result that health types are important is not an artifact of projecting a high-order Markov process into a low-order one.

Second, in all specifications, the estimated coefficients on health types imply that a person with the best (worst) health type has a lower (higher) probability to be in poor health regardless of his current health status.

2.4.1 Predicting health types

Table 2 reports our estimated coefficients for the ordered-logit predicting one's health type which we specify in Eq (6). It shows that one's observed health status h_{t_0}, fixed labor productivity γ, and net worth k_{t_0} are informative about one's health type.

A larger positive coefficient on a variable implies that a person is more likely to be of the worst health type (η_1) and less likely to be of the best health type (η_3). As expected, if an individual's health at age t_0 is good (2^{nd} row), he is significantly less likely to be of η_1-type. The opposite is true if his health is poor (3^{rd} row).

In addition, in all specifications, the coefficients on both fixed productivity γ and net worth k_{t_0} are negative and decreasing in fixed productivity and net worth. Thus, people with the worst health type are more likely to also have low fixed productivity. This can help explain the patterns documented in Section 2.2. Moreover, after controlling for fixed productivity, people from the upper quintile of the wealth distribution are less likely to be of the worst (η_1) health type.

Two forces can give rise to the fact that initial wealth at the first period that we observe people is predictive of their health type conditional on fixed labor productivity. First, the correlation between health type and patience might generate that healthier and more patient people save more. Second, people with bad health types are more likely to be sick for long, to incur large earnings and medical expenses losses, and thus to have lower wealth. It is plausible to think that the second pathway plays a relatively small role at younger ages because typically at those ages medical costs and earning losses due to bad health are small.

To further evaluate the importance of the first pathway, that is the correlation between patience and health types, we re-estimate our health model using only individuals whose initial wealth k_{t_0} is observed at younger ages (that is, age 39 and younger).[5] These results are similar to those that we obtain from the full sample. We report them in Appendix B.4. They confirm that people's wealth the first time we observe them is informative about their

[5] The average initial age of observation in this restricted sample is between 25 and 30, depending on the specification of lag dependence T.

	T=5	T=4	T=3	T=2	T=1
Coefficients of history-dependence terms and health types in Eq. (2)-(3)					
a_2^B	0.0752	0.0715	0.129	0.288	–
a_3^B	0.826***	0.810**	0.675**		
a_4^B	0.528	0.704**			
a_5^B	0.772**				
$a_{\eta_1}^B$	2.270***	1.994***	1.604***	1.987***	2.111***
$a_{\eta_3}^B$	-2.043***	-2.104***	-1.346*	-1.303**	-1.506**
Coefficients of history-dependence terms and health types in Eq. (4)-(5)					
a_2^G	-0.391*	-0.366*	-0.369*	-0.770***	–
a_3^G	-0.242	-0.183	-1.086***		
a_4^G	-1.007***	-1.691***			
a_5^G	-1.921***				
$a_{\eta_1}^G$	4.527***	3.784***	1.637***	1.806***	2.006***
$a_{\eta_3}^G$	-1.447***	-1.639***	-2.318***	-2.555***	-2.871***
N	9028	9765	11126	12096	13083

Table 1: Estimation results of the health shock process in Eq.(2)-(5). $a_{\eta_2}^B$ and $a_{\eta_2}^G$ are normalized to zero. Being in bad/good health for one period ($\tau_B = 1, \tau_G = 1$) is set as the base case. All estimations include a quadratic in age whose coefficients depend on current health status (poor, fair, good).
(* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$)

	T=5	T=4	T=3	T=2	T=1
age t_0	-0.039	-0.025	0.005	0.001	-0.016
$h_{t_0} = G$	-1.457***	-1.429***	-1.879***	-1.921***	-2.250***
$h_{t_0} = P$	1.463	2.072*	2.410	2.386	1.022
2^{nd} tercile of γ	-0.247	-0.337	-0.509**	-0.546**	-0.642***
3^{rd} tercile of γ	-1.203***	-1.374***	-1.188***	-1.286***	-1.355***
2^{nd} quintile of k_{t_0}	-0.002	-0.129	-0.048	-0.459*	-0.469*
3^{rd} quintile of k_{t_0}	-0.620	-0.429	-0.367	-0.378	-0.603**
4^{th} quintile of k_{t_0}	-0.749	-0.606	-0.691*	-0.701**	-0.759***
5^{th} quintile of k_{t_0}	-2.348***	-1.616***	-1.169***	-1.280***	-1.264***

Table 2: Estimation results of the equation predicting health type in Eq.(6). Fair health status, 1^{st} tercile of γ and 1^{st} quintile of k_{t_0} are set as the base for the corresponding dummy variables. All estimations include dummy variables for 10-years bracket of birth year. (* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$)

.

health types and is significant for those in the top wealth quintile, even in this younger sample. This suggests that the correlation between health types and patience is important and further supports our choice of incorporating this mechanism in our life-cycle model.

2.4.2 Transition probabilities

In our life-cycle model, we use our estimated health process with $T = 3$, which means 4 years of lagged dependence. We make this choice to keep our computational costs manageable while preserving the rich dynamics of our health process.

Figure (4) plots the estimated transition probabilities to poor and good health for this version of our health process - the full set of health transition probabilities are shown in Figure B2 in the Appendix. The top, middle, and bottom panels refer to the case in which people are currently in poor, fair and good health, respectively. Different line style represent different health types, while colors refer to different health histories. The key take-aways are as follows.

First, for all health states, as people get older, the probability of moving to a worse health state increases, while the probability of moving to a better health state declines. The speed of the increase/decline varies both by heath type and the duration of the current health state.

Second, one's health status is persistent: the longer an individual spends in the current health state, the higher is the probability to stay in this state. The difference is especially pronounced for those who spend at least three periods in the current health state ($\tau \geq 3$) versus those who only spend two periods ($\tau = 2$). The difference between the history of one ($\tau = 1$) versus two periods ($\tau = 2$) is smaller, especially for people in fair and poor health.

Third, the variation in transition probabilities across health types is much larger than that across health histories. For example, consider a 60-year old person who is in poor health for one period (top left panel). If he is of the best health type η_3 (red line with dots) he has an 20% probability of staying in poor health next period, while if he is of the worst health type η_1 (red solid line), this probability increases to 80%. If, instead, he has been in bad health for three periods, the corresponding probabilities increase approximately by 10 percentage points for both best and worst health types.

The large variation by health types is similar for the transition probabilities from fair and good health. This implies that health types play an important role in explaining the high persistence of one's health status and occurrence of long spells of bad health. In Appendix B.7 we illustrate this point in more details by showing the distribution of people by the number of periods in bad health over the life-cycle.

Consistent with our findings on the importance of health types, Halliday (2008) uses the PSID to estimate a dynamic process for health with fixed heterogeneity and heterogeneous persistence coefficients and finds that, for a large part of his sample, persistence is mostly driven by fixed heterogeneity. Lange and McKee (2012) estimate a dynamic latent health process using multiple health measures from HRS and document that heterogeneity across

individuals (random effects) is important in capturing the high persistence of objective and self-reported health measures.

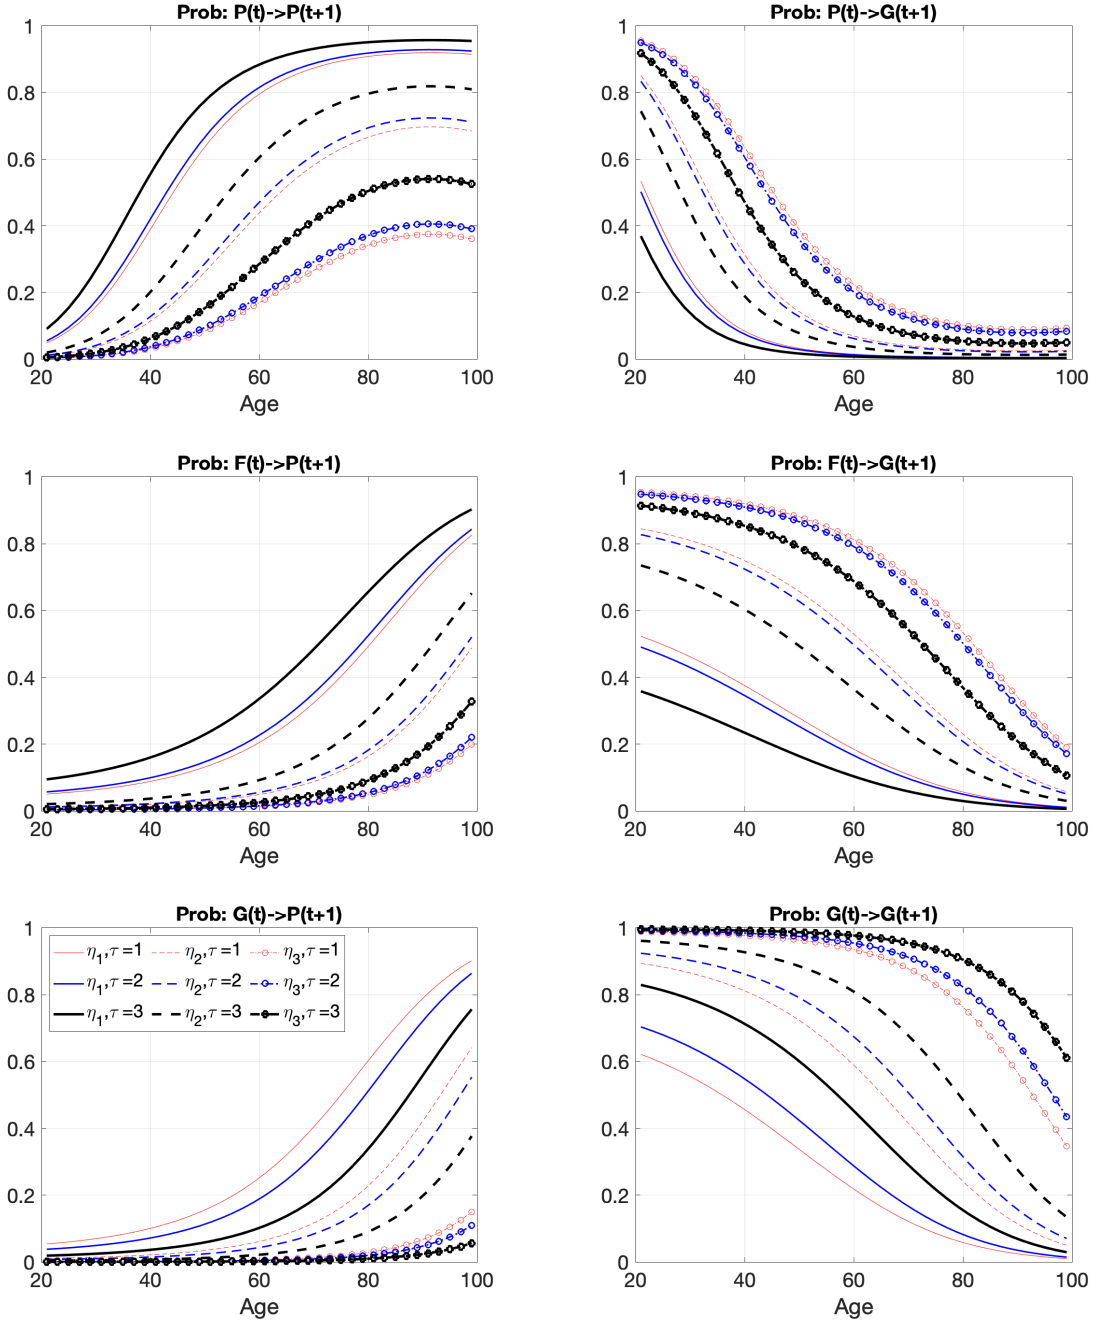


Figure 4: Estimated two-year transition probabilities ($T = 3$). Different line styles represent three health types: solid line for η_1-type, dashed line and line with dots for η_2- and η_3-type, respectively. Black, blue, and red lines correspond to the number of consecutive periods in current health status (τ).

	η_1	η_2	η_3
$Pr(\eta)$	0.08	0.35	0.57
$Pr(\eta \mid \gamma_L)$	0.13	0.44	0.43
$Pr(\eta \mid \gamma_M)$	0.08	0.36	0.56
$Pr(\eta \mid \gamma_H)$	0.04	0.24	0.72

Table 3: Top line: Initial distribution of individuals over health types at age 21. Other lines, joint distribution of $\{\eta, \gamma\}$ at age 21. The terms γ_L, γ_M, and γ_H refer to the bottom, middle, and top terciles of the estimated fixed productivity distribution, respectively.

2.4.3 Initial distributions

Table 3 reports the initial distribution of individuals over health types and the joint distribution over health types and fixed labor productivity, based on our estimates of Eq (6) when $T = 3$. Our ordered-logit model allows the initial distribution of individuals over health types to be asymmetric over three discrete points and to depend on the empirical joint distribution of health status (poor/fair/good) and fixed productivity γ. We use the distribution of health status and estimated fixed productivity γ of people aged between 19 and 24. Almost all individuals in this age range have been either in fair health for one period or in good health for at least three periods. We report the joint distribution of health status and fixed productivity in Table B5 in Appendix B.4.

At age 21, only 8.3% of people are of the worst health type. The rows from 2^{nd} to 4^{th} of Table 3 show that there are proportionately more η_1- and less η_3-types among people in the bottom tercile of the fixed productivity distribution (γ_L) compared to the top tercile (γ_H). In other words, less productive people are also more likely to be of the worst health type.

2.4.4 Comparing the implications of our processes with the data

We now turn to showing that our health and labor productivity processes generate data that are consistent with those from the PSID and the HRS. Figures 1 and 2 show that they match well the fraction of people in bad health by age and the two-year transition probabilities between good and bad health by age, as well as their duration-dependence.

For an additional external validation of our health shock process, we turn to a balanced panel of males with a high-school degree aged 55 to 66 years old in the HRS. Figure 5 reports the number of waves spent in bad health over a ten-year period, conditional on one's initial health at age 55-56. Since the HRS is a bi-annual survey, an individual can only report being unhealthy for at most five periods over the ten year period that we focus on. The darker bars refer to the HRS data, the medium-shaded ones refer to the PSID data, and the lightest ones come from simulations from our estimated processes. The PSID data in the left graph of Figure 5 show that while almost 70% of those initially healthy at age 55-56

continue being healthy in all five waves, a non-trivial fraction (5.7%) becomes unhealthy and stays unhealthy for 4-5 waves during the next ten years. The right panel shows that among those already unhealthy at age 55-56, 51% stay unhealthy during the next 4-5 waves. Overall, there is a substantial fraction of people with very long spells of bad health. This Figure also shows that our estimated processes match the long-term distribution of people in bad health.

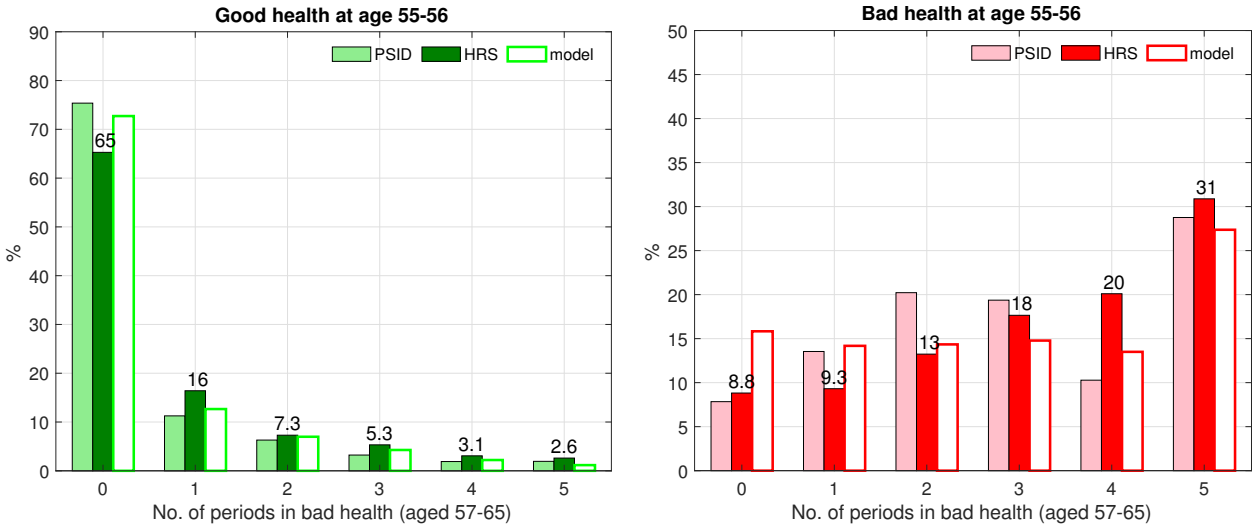


Figure 5: The number of waves being in bad health between age 57-58 and 65-66 conditional on health status at age 55-56. Bad health includes both fair and poor health.

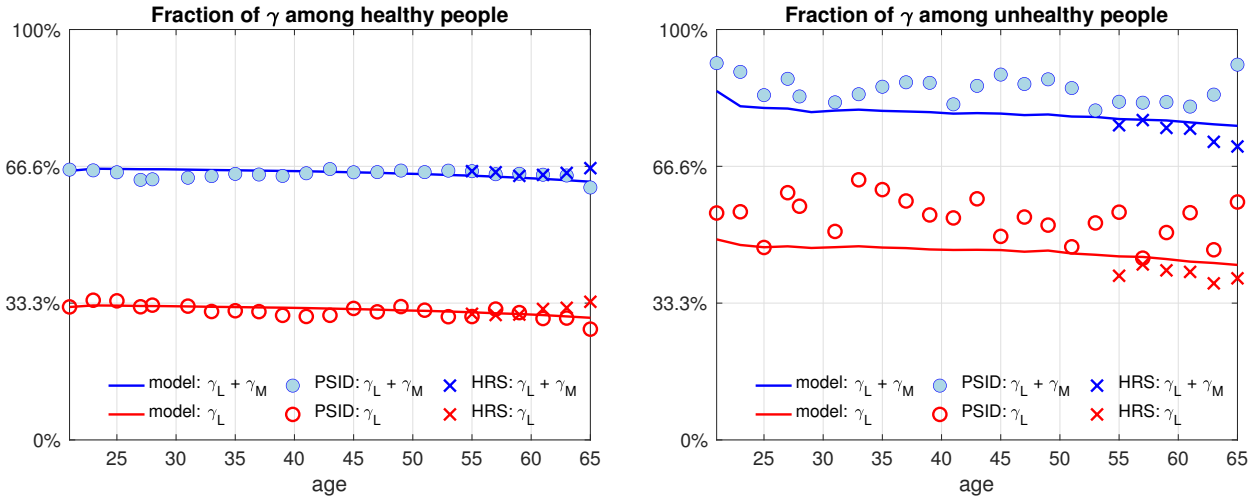


Figure 6: Fraction of people in the bottom two terciles of fixed labor productivity (red lines) or bottom tercile of labor productivity (blue lines) among those in good health (left panel) and bad health (right panel). Bad health includes both fair and poor health. The terms γ_L and γ_M correspond to the lowest and middle terciles of the fixed productivity distribution.

As we have seen in Section 2.2, there are more people with low fixed productivity among

the unhealthy. Hence, as a final validation exercise, we report the joint distribution of health and fixed labor productivity. Figure 6 displays the fraction of people in the bottom two terciles of fixed labor productivity (red lines) or bottom tercile (blue lines) among those in good health (left panel) and bad health (right panel) from the PSID data, the HRS data, and the data simulated from our estimated processes, during all of the working period. The terms γ_L and γ_M correspond to the lowest and middle terciles of the fixed productivity distribution, respectively. This graph reveals that our statistical model for health and productivity also reproduces these features of the data well.

It is important to point out that, in our simulations, we fix the initial joint distribution of health status and fixed productivity at age 21 as we observe it in the data. Hence, the relationship between fixed productivity and health later in the life-cycle is generated solely by two forces. First, the correlation between one's health type (η) and fixed productivity (γ) at age 21, and second by how one's health type affects one's evolution of health. At age 21, the percentage of γ_L-people (worst productivity) among the unhealthy is above 50%. If health type and fixed productivity were independently distributed, after the age of 21 this number would quickly decrease since the fraction of newly sick people of each γ-type will always be 1/3. Note that we do not directly exploit this joint evolution of health and labor productivity in our estimation, i.e., the health transitions are not a function of fixed labor productivity (see Eqs (2)-(5)). Thus, the ability of our model to reproduce it gives our model additional credibility.

2.5 What are health types?

As we have seen, health types play an important role in determining health persistence. They also imply that people differ in their predisposition to fall sick and to recover from illness.

Health types can be related to variation in genetic predisposition and/or lifestyle, where the latter can be partly due to habits developed in childhood. In fact, a number of studies find that childhood circumstances have a long lasting effect on adult health.

For evidence supporting these mechanisms, we turn to the HRS, which contains a wealth of information that is useful to study these questions. From it, we use a balanced panel of healthy men age 55-56, and whom we observe until they are aged 65-66 (the same sample used to construct Figure 5).

Table 4 sorts our HRS sample based on the total number of unhealthy periods that they report over the subsequent ten-year interval. It shows a correlation between the *future* number of unhealthy periods and factors that can be linked to lifestyle *recorded at age 55-56* or genetics.

# unhealthy periods	Individuals' characteristics[a] (HRS)					model	
	% smoking	BMI[b]	% father alive	% mother alive	parents' educ (yrs)[c]	% η_1	% η_3
0-1	22.6	27.9	21.6	48.4	10.1 / 10.5	0.1	78
2-3	27.1	29.5	21.5	50.4	9.2 / 9.9	3	12
4-5	44.4	29.8	16.1	36.5	8.4 / 9.2	25	2

Table 4: Characteristics of a balanced panel of people who are healthy at age 55-56 and alive for the subsequent ten years, organized by the number of unhealthy periods between ages 57 and 66. BMI is the average Body Mass Index. The first and second numbers in the education column refer to average education years of father and mother in each cell, respectively.

# unhealthy periods	Polygenic scores (HRS)			
	educational attainment	smoking	BMI	longevity
0-1	-0.120	0.003	-0.006	-0.06
2-3	-0.216	0.023	0.127	-0.065
4-5	-0.708	0.092	0.140	-0.250

Table 5: Average polygenic scores of people who are healthy at age 55-56 and alive for the subsequent ten years, organized by the number of unhealthy periods between ages 57 and 66.

In particular, individuals who report being unhealthy for four to five periods between ages 57 and 66 are much more likely, at age 55, to have ever smoked, to have a higher body mass index (BMI), and less likely to have living parents. In addition, people with longer unhealthy spells have less educated parents. This is consistent with the findings of Case et al. (2002), who show that parental income and education have a significant impact on child's health and thus on subsequent health in adulthood.

Overall, Table 4 shows that even among a relatively homogeneous sample of healthy males at age 55-56 with the same educational attainment, there is heterogeneity in some fixed or long-lasting factors, which in turn are correlated with their health evolution over the next ten years.

These features of the data are consistent with our model of health dynamics: the last two columns of Table 4 show that in a comparable sample simulated from our model, there are substantial differences in health type composition. While among people with at most one unhealthy period, 78% are of the best health type and almost no one is of the worst health type, among people experiencing 4 to 5 unhealthy periods, only 2% are of the best health type and 25% are of the worst health type.

Table 5 documents the relationship between the number of unhealthy periods (our proxy for health types) and genetic variables that have been found to predict key economic outcomes: polygenic scores. That is, polygenic scores (PGS) are indices created by combining genetic markers that predict certain individuals' outcomes (for more details, see Papageorge and Thom (2018) and Barth et al. (2020)).

We report four polygenic scores (PGS): for educational attainment, for lifestyle-related behaviors such as smoking and BMI, and for longevity. A higher score refers to genetic variation that predicts higher education level, higher propensity to be a smoker, to have higher BMI, and a longer lifespan, respectively. Each scores is normalized to have mean of zero and variance of one. Table 5 reports the average scores of the same HRS sample as the previous table and also sorts by the number of unhealthy periods.

People reporting 4-5 unhealthy periods between the ages of 57 and 66 have, on average, a noticeably lower score for educational attainment and longevity, and higher scores for smoking and BMI compared with those who report at most one unhealthy period. Appendix C.1 shows that this finding is robust when we use the $25^{th}, 50^{th}$, and 75^{th} percentiles of the polygenic scores distribution.

These features of the data confirms the view that health outcomes can, to a certain extent, be traced back to factors determined early in life. Because our agents start adult life at age 21, we capture these features of the data as health types in our stylized model.

In addition, the existence of pre-determined factors affecting health is supported by a growing empirical literature, which we review in Appendix C.2.

3 Our life-cycle model

In this section, we develop a life-cycle model with health uncertainty. In it, health affects people through multiple channels and evolves according to the processes described in the previous section.

3.1 Demographics, preferences, and labor income

A model period is two years long and each individual lives at most T periods. During the first $R-1$ periods of life people chooses whether to work or not. At age R everyone retires. We denote the health-dependent survival probability from age t to $t+1$ as ζ_t^h.

At age t an agent's health, h_t, can be either good (G), fair (F), or poor (P), and next period's health status depends on current health status, the length of the current health spell τ, and health type $\eta_i \in \{\eta_1, \eta_2, \eta_3\}$. See Section 2.3 for more details.

Health and economic outcomes in our model are linked via two mechanisms. First, health directly affects medical spending, productivity, disutility from work, access to health insurance, and survival probabilities. These direct effects have been used in other structural models with health uncertainty, including by Capatina (2015), French (2005), French and Jones (2011), Pashchenko and Porapakkarm (2013, 2016a, 2017), and Rust and Phelan (1997).

Second, individuals differ in fixed ex-ante characteristics, that is in their fixed labor productivity, health type, and patience, and these characteristics can be correlated with each other.[6] This correlation captures the fact that some factors that are determined early in life can influence all three fixed factors (see discussion in Section 2.4.1). This mechanism creates an additional compositional difference between the healthy and the unhealthy. That is, among the unhealthy there can be more impatient people, and this can partially account for the observed disparities in economic outcomes between the two groups.

Formally, we assume that the discount factor (β_i) can take two values, $\beta_i \in \{\beta_{low}, \beta_{high}\}$, where $\beta_{low} < \beta_{high}$. At age 21 (when an individual enters the model) the joint distribution of the discount factor and health type, $\{\beta_i, \eta_i\}$, is captured by $Pr(\beta_j|\eta_m) \in [0,1]$, where $j \in \{low, high\}$ and $m \in \{1,2,3\}$. Since, in turn, the health type can be correlated with one's fixed labor productivity, all three fixed factors can be correlated with each other.

An individual is endowed with one unit of time that can be used for either leisure or work. Labor supply (l_t) is thus indivisible; $l_t \in \{0,1\}$. Work implies a fixed utility cost ϕ_W. For people with fair and poor health, there is an additional disutility from working, denoted ϕ_F and ϕ_P, respectively. We assume that the preferences of individuals over consumption and leisure take the following form:

$$u(c_t, l_t, h_t) = \frac{(c_t/\overline{n}_t)^{1-\rho}}{1-\rho} - \phi_W \mathbf{1}_{\{l_t>0\}} - \phi_F \mathbf{1}_{\{h_t=F, l_t>0\}} - \phi_P \mathbf{1}_{\{h_t=P, l_t>0\}} + \overline{b}, \tag{8}$$

where ρ is risk-aversion and $\overline{n}_t$ is an age-specific household size.[7] We follow Hall and Jones (2007) by adding a positive term $\overline{b}$ to ensure that individuals in our model value their life; i.e., the continuation value of being alive exceeds the utility when deceased. This matters when we compute the welfare costs of bad health, because otherwise sick people would be happy about dying sooner.

As in De Nardi (2004), individuals also derive utility from leaving a bequest of size k

$$\upsilon(k) = \theta_{Beq} \frac{(k + k_{Beq})^{1-\rho}}{1-\rho},$$

[6] Many studies find heterogeneity in patience, including Epper et al. (2018), Lawrance (1991), Warner and Pleeter (2001). In addition, Cronqvist and Siegel (2015) find that genetic differences explain a significant fraction of the variation in saving propensities across individuals, potentially through a link to the rate of time preferences or self-control. A number of macroeconomic studies point out the importance of heterogeneity in rates of time preferences to explain wealth inequality (Carroll et al., 2016, Hendricks, 2007a, 2007b; Krueger et al., 2016; Krusell and Smith, 1998, Samwick, 1998).

[7] We incorporate the family size $\overline{n}_t$ into the model in order to make the wealth profile produced by the model consistent with the profiles constructed from the PSID. (See Section 4.1 for details on the construction of the targeted wealth profile). We take the average family size $\overline{n}_t$ from the PSID. The average family size is 1.9 at age 21, it increase to 3.18 at age 39 before declining to 2.07 at age 65 and further to 1.02 at age 99.

where θ_{Beq} determines the strength of the bequest motive and k_{Beq} determines to what extent bequests are a luxury good.

Earnings are given by $z_t^h l_t$, where z_t^h is an idiosyncratic productivity component given by

$$z_t^h = \lambda_t^h \Upsilon_t, \tag{9}$$

where λ_t^h is a deterministic function of age and current health (good, fair or poor), while Υ_t is the stochastic shock that we specify in Section 3.4.

3.2 Medical expenses and health insurance

During each period, every agent receives a medical expense shock $\left(x_t^h\right)$ which depends on age and current health. We denote the distribution of medical shocks as $\mathcal{G}_t(x_t^h|h_t)$. Appendix D.1 discusses how we estimate these shocks.

Individuals in our model also differ in their health insurance status: working age people can be uninsured or covered by different types of private insurance, that is individual or employer-sponsored health insurance (ESHI). In contrast, retirees are covered by Medicare. We index the individual's insurance status by using i_H, where $i_H = 0$ corresponds to being uninsured, $i_H = 1$ corresponds to having individual insurance, $i_H = 2$ corresponds to group insurance (or ESHI), and $i_H = 3$ corresponds to Medicare. All types of insurance only provide partial medical expenses coverage. We denote by $cvg\left(x_t^h, i_H\right)$ the fraction of medical expenses covered by insurance and allow it to be a function of one's medical shock and insurance type. Note that $cvg\left(x_t^h, 0\right) = 0$.

A working-age individual receives an offer to buy ESHI with probability $Prob_t$, which depends on age (t), fixed labor productivity (γ), and health ($h_t \in \{G, F, P\}$). We estimate $Prob_t$ from the MEPS. The variable g_t characterizes the status of the offer: $g_t = 1$ if an individual gets an offer, and $g_t = 0$ otherwise. Only working individuals with an offer ($l_t = 1$, $g_t = 1$) can purchase the ESHI insurance. We assume that an employer who offers ESHI fully covers the premium, i.e., the employer contribution is 100%. On average, employers who offer ESHI contribute about 80% of the premium for single coverage and about 70% for family coverage (Kaiser Family Foundation, 2004). We abstract from workers' contribution for simplicity. This assumption does not affect our results but helps lower computational costs because working individuals with an ESHI offer always buy insurance.

Every working-age individual can buy health insurance in the individual health insurance market at the price $p_I\left(h_t, t\right)$, which depends on one's age and health. We assume that an individual's insurance premium is based on his expected medical costs and administrative

loads:

$$p_I\left(h_t,t\right)=\xi EM_t\left(h_t,t\right)+\varphi^h. \tag{10}$$

The term ξ is a proportional load, while φ^h is a fixed load. We allow the fixed load to depend on health to capture the fact that unhealthy individuals may face more frictions when purchasing insurance through the individual market, for example, through search costs or a larger probability of being denied coverage due to pre-existing conditions.

The expected medical costs covered by insurance are

$$EM_t\left(h_t,t\right)=\sum_{x_t^h}x_t^h cvg\left(x_t^h,1\right)\mathcal{G}_t(x_t^h|h_t).$$

We denote the Medicare premium as P_{MCR}. This corresponds to the Medicare Part B premium.

3.3 Taxation and social transfers

We model the tax system as follows. Working households pay payroll taxes, which include the Medicare tax (τ_{MCR}) and the Social Security tax (τ_{ss}). The latter only affects earnings below $\overline{y}_{ss}$. There is a consumption tax τ_c and a tax on capital income τ_k. There is a progressive labor income tax $\mathcal{T}\left(y\right)$ which we specify as Heathcole et al. (2020)

$$\mathcal{T}\left(y\right)=y-a_{\tau 0}\ y^{1-a_{\tau 1}}. \tag{11}$$

The progressivity of the tax system is captured by $a_{\tau 1}$. We explain how we set taxes and transfers in Appendix D.2.

We represent several existing means-tested programs (Medicaid, food stamps, Disability Insurance, and Supplement Security Income) in a stylized way through a public safety-net program, $T^{SI}(\overline{c})$. This program guarantees every household a minimum consumption floor $\overline{c}$. This floor also captures the existence of uncompensated care or medical bankruptcy. In fact, in 2004, 85 percent of the uncompensated care was paid by the government.

Retirees receive Social Security benefits ss. In practice, these payments depend on an individual's history of earnings. To capture the existing variation in pension benefits without increasing computational costs, we approximate the benefits using the following approach. First, we divide individuals into groups based on their health just before retirement h_{R-1}, on their last draw of the persistent productivity shock ν_{R-1} (see Section 3.4), and on their fixed ex-ante heterogeneity (γ,η,β). Then, for each group, we compute average earnings over the 17 model periods (34 years) with the highest earnings. Then we apply the Social Security

benefits formula to these average earnings. This way, Social Security benefits in our model can be represented as $ss(h_{R-1}, \nu_{R-1}, \gamma, \eta, \beta)$.

3.4 The labor productivity shock

Our labor productivity process is given by

$$z_t^h = \lambda_t^h \Upsilon_t = \lambda_t^h \exp(\nu_t) \exp(\gamma), \tag{12}$$

$$\begin{aligned} \nu_t &= \rho_\nu \nu_{t-1} + \varepsilon_t; \quad \varepsilon_t \sim N(0, \sigma_\varepsilon^2), \\ \gamma &\sim N(0, \sigma_\gamma^2), \end{aligned} \tag{13}$$

where λ_t^h is a deterministic component that depends on age and health ($h_t \in \{G, F, P\}$). The idiosyncratic component Υ_t consists of a persistent shock ν_t and fixed productivity γ.

We detail our estimation of the stochastic part of productivity in Appendix D.3. Following French (2005), we estimate the deterministic component of labor productivity λ_t^h at the second stage of our estimation. We do this because the fraction of unhealthy workers is significantly below 100%, so the average income conditional on working could be a biased estimate of λ_t^h if there is selection into employment.[8]

3.5 Timing of the model

The timing of the model is as follows. At the beginning of the period, individuals learn their productivity, health and ESHI offer status. Based on this information, an individual decides his labor supply (l_t) and insurance choice (i_H). At the end of the period, the medical expenses shock (x_t^h) is realized. After paying the out-of-pocket medical expenses, an individual chooses his consumption (c_t) and savings for the next period (k_{t+1}). The problem of retirees is simpler; they only choose consumption and savings for the next period.

3.6 The optimization problem

Working age individuals ($t < R$)**.** At the beginning of each period, the state variables for an individual i are capital (k_t), health status ($h_t \in \{G, F, P\}$), length of the current health spell ($\tau \in \{1, 2, 3\}$), productivity shock (ν_t), ESHI offer status $\big(g_t \in \{0, 1\}\big)$, age ($t \in \{1, 2, ..., R-1\}$), fixed productivity ($\gamma \in \{\gamma_L, \gamma_M, \gamma_H\}$), health type ($\eta \in \{\eta_1, \eta_2, \eta_3\}$)

[8] An alternative method is to apply the two-step Heckman estimation directly to the data. However, this approach requires a valid variable serving as an exclusion restriction. Though more computationally costly, our strategy ensures that the model reproduces the income-health gradient in the data (Figure 9 and Table 7), which is important for evaluating the costs of bad health through the labor market channel.

and discount factor ($\beta \in \{\beta_{low}, \beta_{high}\}$). To make our expression less cluttered, we omit the subscript i for all state variables.

We denote the vector of state variables as $\mathbb{S}_t$.

The value function of a working age individual at the beginning of period t is:

$$V_t\left(\mathbb{S}_t\right) = \max_{l_t, i_H} \sum_{x_t^h} \mathcal{G}_t(x_t^h|h_t) W_t(\mathbb{S}_t; l_t, i_H, x_t^h) \tag{14}$$

where

$$W_t(\mathbb{S}_t; l_t, i_H, x_t^h) = \max_{c_t, k_{t+1}} u\left(c_t, l_t, h_t\right) + \beta \left[\zeta_t^h E_t\Big(V_{t+1}\left(\mathbb{S}_{t+1}\right)\Big) + \left(1-\zeta_t^h\right)\theta_{Beq}\frac{(k_{t+1}+k_{Beq})^{1-\rho}}{1-\rho}\right] \tag{15}$$

subject to

$$k_t\left(1+\left(1-\tau_k\right)r\right) + z_t^h\, l_t - x_t^h\left(1-cvg(x_t^h, i_H)\right) - P_t^h - Tax + T^{SI}\left(\overline{c}\right) = (1+\tau_c)c_t + k_{t+1} \tag{16}$$

$$P_t^h = \begin{cases} 0 & ;\text{ if } i_H \in \{0,2\} \\ p_I\left(h_t, t\right) & ;\text{ if } i_H \in \{1\} \end{cases} \tag{17}$$

$$T^{SI}(\overline{c}) = max\Big(0, (1+\tau_c)\overline{c} + Tax + P_t^h + x_t^h\big(1-cvg(x_t^h, i_H)\big) - k_t(1+(1-\tau_k)\,r) - z_t^h l_t\Big)$$

$$Tax = \mathcal{T}\left(z_t^h l_t\right) + \tau_{MCR} z_t^h l_t + \tau_{ss}\min\big(z_t^h l_t, \overline{y}_{ss}\big) \tag{18}$$

$W_t(\mathbb{S}_t; l_t, i_H, x_t^h)$ is the interim value function conditional on the labor supply and insurance choices after the medical shock is realized. The conditional expectation on the right-hand side of Equation (15) is over $\left\{h_{t+1}, z_{t+1}^h, g_{t+1}\right\}$. Equation (16) is the budget constraint; in this constraint P_t^h is the insurance premium, which is described in Equation (17). In Equation (18), the first term is the income tax and the last two terms are payroll taxes.

Retired individuals ($t \geq R$). The state variables for retired people are assets (k_t), health status (h_t), length of the current health spell (τ), medical shock $\left(x_t^h\right)$, health status at 64 ($h_{R-1} \in \{P, F, G\}$), productivity shock before retirement (ν_{R-1}), age ($t \in= \{R, ..., T\}$), fixed producivity type (γ), health type (η), and discount factor (β). We denote the vector of state variables as $\mathbb{S}_t^R$.

The value function of a retired household is:

$$V_t\left(\mathbb{S}_t^R\right) = \sum_{x_t^h} \mathcal{G}_t(x_t^h|h_t) W_t(\mathbb{S}_t^R; x_t^h) \tag{19}$$

where

$$W_t(\mathbb{S}_t^R; x_t^h) = \max_{c_t, k_{t+1}} u\left(c_t, 0, h_t\right) + \beta \left[\zeta_t^h E_t\Big(V_{t+1}\left(\mathbb{S}_{t+1}^R\right)\Big) + \left(1-\zeta_t^h\right)\theta_{Beq}\frac{(k_{t+1}+k_{Beq})^{1-\rho}}{1-\rho}\right] \quad (20)$$

subject to

$$k_t\left(1+(1-\tau_k)\,r\right)+ss-x_t^h\left(1-cvg(x_t^h,3)\right)-P_{MCR}-\mathcal{T}\left(ss\right)+T^{SI}\left(\overline{c}\right)=(1+\tau_c)c_t+k_{t+1} \quad (21)$$

$$T^{SI}(\overline{c}) = max(0, (1+\tau_c)\overline{c} + \mathcal{T}\left(ss\right) + P_{MCR} + x_t^h(1-cvg(x_t^h,3)) - k_t(1+(1-\tau_k)\,r) - ss)$$

$W_t(\mathbb{S}_t^R; x_t^h)$ is the interim value function conditional on medical shock realization. The conditional expectation on the right-hand side of Equation (20) is over h_{t+1}. Equation (21) is the budget constraint.

4 Model estimation

In this section, we explain our strategy to estimate the model parameters, describe the estimation results, and illustrate the fit of the model to the data, including non-targeted moments.

4.1 Estimation strategy

We adopt a two-step estimation strategy. In the first step, we set parameters related to demographics, taxes, social security benefits, and health insurance, and estimate the shock processes directly from the data. We explain how we estimate our first-step parameters in Appendix D. The survival probability and health shock process are taken from Section 2.4. We fix the interest rate r at 2%. We set risk aversion ρ to 3, a value commonly used in structural life-cycle studies.

Given the parameters and the shock processes from the first step, we implement the Method of Simulated Moments to estimate our remaining model parameters. We minimize the weighted sum square of difference between the targeted and simulated moments using the inverse of squared standard errors as the weights. The set of parameters estimated at the second stage is $\{\overline{b}, \phi_W, \phi_F, \phi_P, \lambda_t^G, \lambda_t^F, \lambda_t^P, \beta_{low}, \beta_{high}, Pr\left(\beta_{low}|\eta\right), Pr\left(\beta_{high}|\eta\right), \theta_{Beq}, k_{Beq}, \overline{c}\}$. Our targeted moments are described below.

The Value of Statistical Life (VSL) We set the target average VSL among the working-age population in our model to \$2 millions. The VSL represents the monetary value corre-

sponding to the reduction in mortality risk that would prevent one statistical death. More formally, it is the marginal rate of substitution between wealth and survival probability. In our framework, it can be expressed as follows:[9]

$$VSL_t = \frac{\partial V_t / \partial \zeta_t^h}{\partial V_t / \partial k_t}$$

Viscusi (1993) provided an extensive review documenting that the estimates vary from \$1 million to \$16 millions (in 1990 dollars). The US government agencies (Department of Transportation, Food and Drug Administration, Environmental Protection Agency) use the VSL of \$1-10 millions in their analysis involving a mortality risk (Robinson, 2007). Because we set the targeted VSL to the lower end of the empirical estimates, we obtain a lower bound for the non-pecuniary effects of health. As discussed below, our results show that even under this parameterization the non-pecuniary effects of health are large. In Appendix E.3 we report the results from an alternative parameterization when the VSL is equal to \$6 millions, which emphasizes the non-pecuniary effects of health more.

Labor market outcomes We target the fraction of workers and average labor income conditional on working for each age and health status (good, fair, and poor). Figure 9 reports our targeted moments from the PSID as dots, our model implications as solid lines, and the HRS data as crosses. The labor income profiles in the right panel come from our estimated coefficients in Eq (1) and include fixed productivity (γ) for each age and health. Since the estimated labor income profiles among workers in poor and fair health are quite similar, we estimate their targeted labor income profiles from the pooled sample of people in poor and fair health. On this point, it is worth pointing out that, even though the targeted labor income profile for people in poor and fair health is the same, their estimated deterministic productivities (λ_t^F, λ_t^P) are different due to their selection into employment.

Wealth moments We target the $25^{th}, 50^{th}$, and 75^{th} percentiles of wealth, conditional on being healthy and unhealthy (poor+fair) by 5-year age brackets (dashed lines in Figure 7). We discard the wealth moments below age 25 because we assume that individuals enter the model with zero assets.

To construct our targeted wealth profiles, we use net worth from the PSID (1994, 1999-2017).[10] Because net worth is measured at the household level and our model abstracts from

[9] Since $\bar{b}$ enters the utility function additively, we can estimate $\bar{b}$ that reproduces the targeted VSL after getting the estimates of the other parameters.

[10] Net worth is given by the sum of the values of business/farm, checking/saving accounts, real estate, stock, vehicles, other assets, annuity/IRA accounts, and home equity, less the value of mortgages/debts. We

heterogeneity in family size, we adjust observed wealth by family size as follows:

$$wealth_{it} = \sum_{j=G,B} \left(d^j_{age} D^{age}_{it} + d^j_1 n_{it} + d^j_2 n^2_{it} + d^j_3 n^3_{it} \right) \mathbf{1}_{\{h_{it}=j\}} + \sum_{t=1994}^{2013} d_t D_t + res_{it}, \qquad (22)$$

where $wealth_{it}$ is net worth, D^{age}_{it} and D_t are age and year dummy variables, and n_{it} is the number of individuals in a family unit. Given the estimated coefficients and the residuals res_{it}, we replace n_{it} in the above equation with the average family size at each age, $\overline{n}_t$, to get our measure of net worth.[11] Then we construct the targeted $25^{th}, 50^{th}$, and 75^{th} percentiles of wealth distribution among people in good and bad health, and report them as dashed lines in Figure 7. As a comparison, we also apply the same method to net worth in the HRS (1994-2012) and plot the results as dotted/crossed marks in Figure 7.

The wealth profiles from the two datasets are remarkably similar. Figure 7 displays the wealth-health gradient typically documented in the literature: Figure 8 emphasizes this gradient by plotting the gap in wealth between the healthy and the unhealthy. This gap starts at relatively young ages and widens until retirement age. This feature of the data suggests that it is important to model the entire life-cycle to understand the costs of bad health.

4.2 Second step estimation results

The third column of Table 6 reports our estimated preference parameters and consumption floor. The discount factors play an important role in wealth accumulation before retirement and its distribution; our estimated β_{low} and β_{high} are 0.877 and 0.992, respectively.[12] The correlation between the discount factor and health type is identified by matching the wealth levels of the healthy and the unhealthy. We find a strong correlation between one's discount factor and health type: the fraction of impatient people (β_{low}) among those with η_1- and η_2-types is about 80%, while the fraction among those with the best health type (η_3) is slightly less than 40%. The average discount factors among η_1- and η_2-types is 0.90 and the average among η_3-type is 0.95. The unconditional average of the discount factor in our model, $E(\beta)$, is 0.932.

The estimated bequest parameters θ_{Beq} and k_{Beq}, which mostly affect wealth decumula-

convert it to 2013 dollars using the CPI.

[11] We compute $\sum_{j=G,B} \left(\widehat{d}^j_{age} D^{age}_{it} + \widehat{d}^j_1 \overline{n}_t + \widehat{d}^j_2 \overline{n}^2_{age} + \widehat{d}^j_3 \overline{n}^3_{age} \right) \mathbf{1}_{\{h_{it}=j\}} + \widehat{d}_{2013} + \widehat{res}_{it}$, where $\widehat{d}^j$ and $\widehat{res}_{it}$ are the estimated coefficients and the residuals from Equation (22). By construction, we remove the variation in net worth due to the variation in family size that is orthogonal to health status and age.

[12] In Appendix E.1 we explain in detail how the difference in estimated discount factors in our model compares to other studies that allow for heterogeneity in the rate of time preferences.

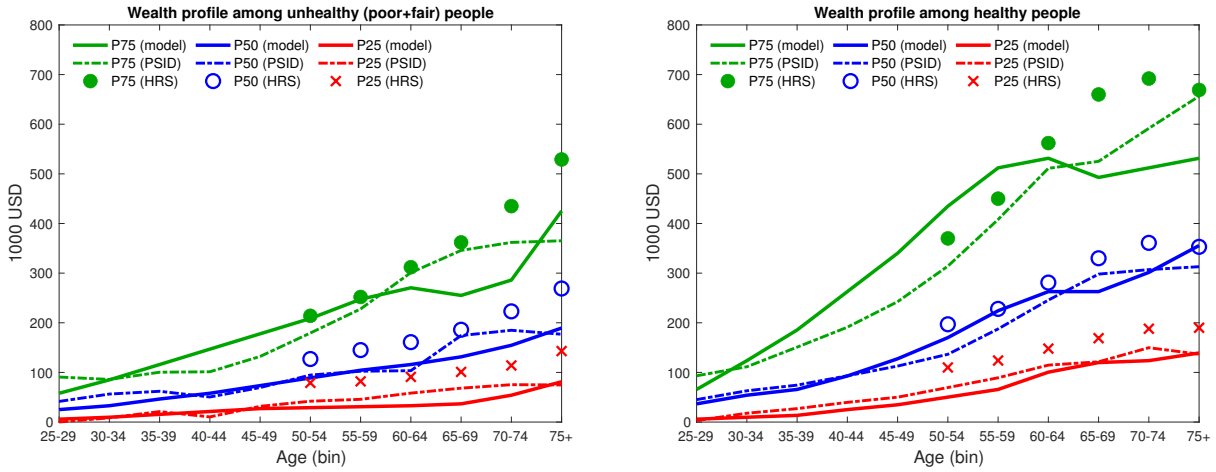


Figure 7: Wealth profiles by health status : data vs model.

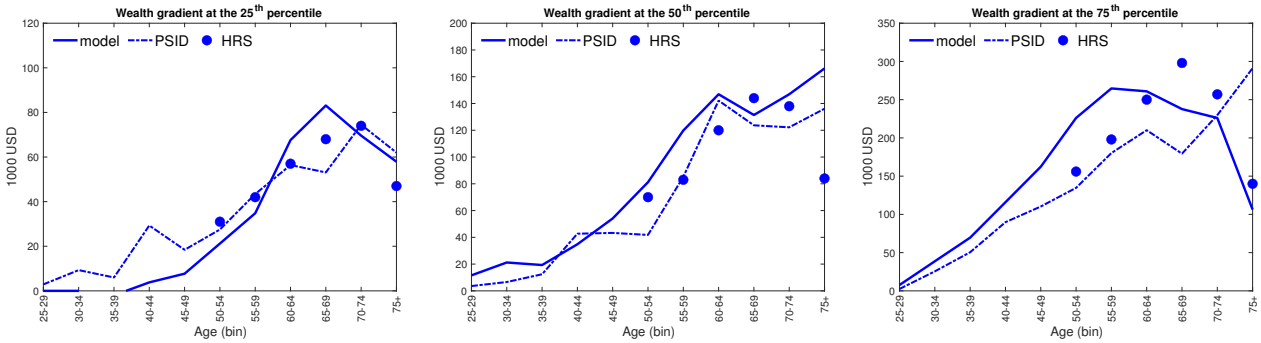


Figure 8: Wealth gradient : data vs model.

Parameters		Baseline	No correlation
(annual) discount factor	$\{\beta_{low}, \beta_{high}\}$	$\{0.877, 0.992\}$	$\{0.88, 0.984\}$
% β_{low} by η_i at age 20	$Pr\left(\beta_{low}\vert\eta_1\right)$	78%	50%
”	$Pr\left(\beta_{low}\vert\eta_2\right)$	79%	50%
”	$Pr\left(\beta_{low}\vert\eta_3\right)$	38%	50%
bequest parameter	θ_{Beq}	1,905	1,787
”	κ_{Beq}	\$182,707	\$199,279
(annual) consumption floor	$\bar{c}$	\$3,505	\$3,373

Table 6: Preference parameters and the consumption floor. Both β and $\bar{c}$ are convented into annual values.

tion after retirement, are 1,905 and 182,707, respectively. In a one-period consumption-saving model with a risk aversion of 3, these values imply that, during the last period of life, the bequest motive becomes operational at an asset level of \$15,000 and the marginal propensity to bequeath (MPB) is 0.92. In other words, individuals with assets below \$15,000 would not

leave bequests, while individuals with assets above $15,000 would leave 92 cents out of every additional dollar for bequests. These numbers are within the range of values found in other studies. For example, the estimation in De Nardi et al. (2010) implies a bequest threshold of about $36,000 and a MPB of 0.88. Pashchenko (2013) provides a comparison of the MPBs and bequest thresholds across several structural life-cycle studies.

The annual consumption floor, which mostly affects the savings of those with lower income, is $3,505. This estimate is consistent with those from other structural models featuring the full life-cycle, medical spending uncertainty, and endogenous labor supply. More specifically, Capatina's (2015) estimate of the consumption floor is $4,114 (in 2006 USD) while Pashchenko and Porapakkarm's (2017) estimate is $1,540 (in 2003 USD).

4.3 Model fit

Figure 9 compares the employment rate (left panel) and the average labor income of workers (right panel) generated by our model (solid lines) with the targeted profiles from the PSID (dots and crosses).[13] Our model matches the important differences in labor market outcomes across health status very well.

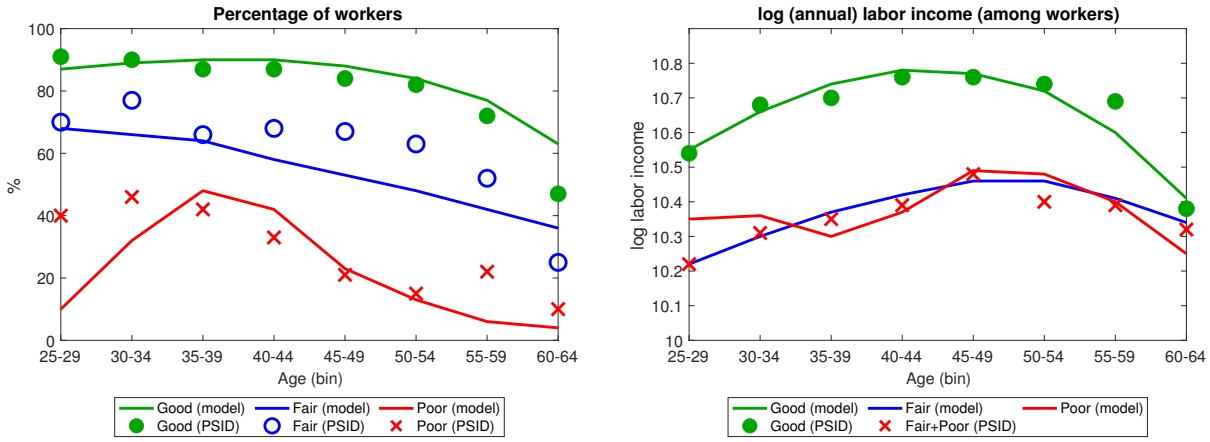


Figure 9: Employment by health (left panel) and average labor income among workers by health (right panel). The dotted and crossed marks are from the PSID while the solid lines are from our model.

Figure 7 displays the wealth profiles from our model (solid line) and the data (with dashed lines for the PSID and dots/crosses for the HRS).[14] Our model matches the wealth

[13] The annual labor income profiles from the model are constructed from dividing two-year labor income by two.

[14] While we capture overall wealth inequality quite well, our model does not generate enough wealth inequality at the very top. Previous literature shows that two important economics mechanisms are crucial to explain that. They are entrepreneurial choices (Quadrini, 1999 and Cagetti and De Nardi 2006) and the intergenerational transmission of bequests and human capital (De Nardi, 2004). For tractability, and because our paper is not about modeling the saving decisions of the very rich, we abstract from these mechanisms.

gap between healthy and unhealthy people for the $25^{th}, 50^{th}$, and 75^{th} percentiles. It is especially worth noting that, even though everyone starts with zero assets in our model, our simulated profiles track very well the widening wealth-health gradient by age. Since the monetary costs of bad health (low earnings and high medical spending) among the young are relatively small compared to older groups, the wealth gradient for the younger group is mostly explained by the larger fraction of β_{low}-individuals among the unhealthy.

Our model also matches additional dimensions of the data by health that we do not target. The first three columns of Tables 7 and 8 show that it replicates the distribution of people by health conditional on both income and wealth. More specifically, the table reports the percentage of unhealthy people by income and wealth terciles. In the data, conditioning on age group, there are much more unhealthy people in the lowest terciles of earnings and wealth. Our model matches these additional features of the data well. Capturing these aspects of the data is important to properly evaluate the long-term effects of bad health.

	PSID (HRS)			Model		
	bottom 1/3	middle 1/3	top 1/3	bottom 1/3	middle 1/3	top 1/3
25-34	14%	5%	3%	11%	3%	1%
35-44	18%	7%	3%	17%	3%	3%
45-54	26%	9%	6%	23%	8%	4%
55-64	30% (38%)	17% (21%)	9% (13%)	32%	10%	7%

Table 7: Percentage of unhealthy (poor+fair) individuals in each earnings tercile. Left panel: data from PSID (data from HRS is in the bracket). Right panel: model.

	PSID (HRS)			Model		
	bottom 1/3	middle 1/3	top 1/3	bottom 1/3	middle 1/3	top 1/3
25-34	10%	10%	6%	6%	5%	4%
35-44	14%	14%	5%	9%	9%	5%
45-54	20%	15%	9%	15%	13%	6%
55-64	32% (37%)	17% (23%)	12% (15%)	23%	18%	8%
65-74	35% (39%)	26% (24%)	16% (16%)	32%	24%	11%
75+	46% (43%)	34% (31%)	21% (28%)	38%	31%	21%

Table 8: Percentage of unhealthy (poor+fair) individuals in each wealth tercile. Left panel: data from PSID (data from HRS is in the bracket). Right panel: model.

5 Results

In this section, we use our estimated life-cycle model to deliver several interesting results. First, we illustrate the importance of the correlation between our ex-ante fixed characteristics in generating the observed health-related wealth inequality. Second, we construct a comprehensive measure of the monetary costs of bad health over the life-cycle. Third, we

evaluate the welfare losses due to bad health realizations, a metric that takes into account both the pecuniary and non-pecuniary consequences of bad health. For both monetary and welfare losses, we first report the average annual losses over the life-cycle and we then provide a decomposition analysis to understand how different aspects of bad health contribute to its total effect.

5.1 Compositional differences and the wealth-health gradient

As we have seen in the previous subsection, our estimates imply a non-trivial compositional difference between the healthy and unhealthy, which is due to the estimated correlation between health types and the rate of time preferences. To illustrate the importance of this compositional difference, we re-estimate an alternative model in which one's discount factor is orthogonal to one's health type and, consequently, to health status. Formally, we restrict $Pr(\beta_j|\eta_m) = 0.5$ for all m and j. We call it the "no-correlation" model. The "no-correlation" model still features all of the channels through which bad health can affect individuals' savings.

In contrast to the baseline model where we use $Pr(\beta_{low}|\eta)$ to capture the wealth-health gradient, the "no-correlation" model has no specific parameters to capture the wealth profiles of the healthy and unhealthy.[15] Because of this, in our estimation, we replace our wealth moments by health with the *unconditional* wealth quartiles; that is, the $25^{th}, 50^{th}$, and 75^{th} percentile of wealth for each 5-year age bracket. It is worth pointing out that even the "no correlation" model requires heterogeneity in time preferences to match the age-profile of the wealth distribution. This is a commonly used approach in the literature that aims at matching wealth inequality (e.g., Hendricks, 2007).

Our estimated parameters for the "no-correlation" model are similar to those from our baseline model, including for the rates of time preferences $(\beta_{low}, \beta_{high})$ (See the last column of Table 6 for their values). In addition, the model-simulated data matches the targeted wealth quartiles unconditional on health status. See the second and fourth columns of Table 9 for a comparison of the $25^{th}, 50^{th}$, and 75^{th} wealth percentiles for the 60-64 age group in the PSID (and HRS) and in the "no-correlation" model, respectively. As a reference, we also report the corresponding statistics from the baseline model. The "no-correlation" model also matches the employment rate and the average labor income by health status conditional on working (the income-health gradient).

[15] In the no-correlation model, wealth is different by health because healthy and unhealthy people have different income, medical spending and life expectancy. However, the difference in income, medical spending, and life expectancy by health is either estimated outside the model or inside the model without taking the wealth-health gradient into account. Because of this, we do not have any parameters left to target the wealth-health gradient once we shut down the correlation between β and η.

Wealth percentile	PSID (HRS)	Baseline	No correlation $Pr(\beta_{low}\|\eta_i) = 0.5$
25^{th} pct	\$92 (\$117)	\$89	\$81
50^{th} pct	\$213 (\$220)	\$224	\$208
75^{th} pct	\$476 (\$383)	\$502	\$445

Table 9: *Unconditional* wealth quartiles at age 60-64 (in 1,000 USD).

Wealth difference by health status	PSID (HRS)	Baseline	No correlation $Pr(\beta_{low}\|\eta_i) = 0.5$
25^{th} pct	\$56 (\$47)	\$67	\$38
50^{th} pct	\$142 (\$98)	\$146	\$38
75^{th} pct	\$210 (\$222)	\$260	\$91

Table 10: Wealth-health gradient at age 60-64 (in 1,000 USD) . The table reports the wealth difference between healthy and unhealthy (poor+fair) people for each wealth quartile.

The "no-correlation" model, however, falls short of replicating the observed large differences in wealth by health status (See Table 10). For example, the difference between the median wealth of the healthy and unhealthy near retirement is only \$38k in this no-correlation model, compared to about \$142k in the data and in the baseline model.

From these findings, we conclude that, even for a relatively homogeneous group of males with the same education level, the direct effect of bad health (low earnings, high out-of-pocket medical expenses and shorter life expectancy) only partially accounts for the observed difference in accumulated wealth between the healthy and unhealthy, and that the income-health gradient does not imply the wealth-health gradient.

Consistent with our findings, Poterba et al. (2017) use HRS data to document that there is a large difference in asset growth between those in the top and bottom one-third of health status (they construct a continuous health index) between the ages of 51 and 61. They also find that only 20-40% of the differences in asset growth can be attributed to the lower earnings and annuity income of those in poor health.

5.2 The monetary losses due to bad health

In our model, the monetary costs of bad health arise from exogenous health-dependent stochastic processes (productivity, medical shock, and ESHI offer probability) and from behavioral responses (labor supply, insurance purchase, and saving decisions). By adding up these costs, we can measure the accumulated losses due to bad health that an individual experiences over the life-cycle.

To understand how costly it is for people to be unhealthy over their life course, we measure how much better off they would be in a counterfactual situation where their health

is always good but everything else is the same. Our structural framework allows us to construct this counterfactual situation. In particular, we consider the following experiment. We place people in an environment where they unexpectedly draw good health realizations every period (they are exceptionally lucky), while the realizations of their other shocks are the same as in the baseline.

This approach has an important advantage over estimating the costs of bad health directly from the data. In the latter case, we can only compare *different* individuals, healthy and unhealthy ones, after controlling for observable characteristics. But, as we showed earlier, healthy and unhealthy people differ in unobservable characteristics (health type, permanent productivity, and preferences) and this biases the empirical estimates of the costs of bad health. With our approach, instead, we compare the exact same people.

To formalize our approach, denote income net of total medical spending of an individual i at time t in the baseline and counterfactual cases as y_{it}^{BS} and y_{it}^{H}, respectively. The difference between y_{it}^{BS} and y_{it}^{H} represents the pecuniary costs of bad health in period t. Our measure of the lifetime costs of bad health averages these costs over the life-cycle and is computed as $\frac{1}{\widehat{T}} \sum_{t=1}^{\widehat{T}} \frac{y_{it}^{H} - y_{it}^{BS}}{(1+r)^t}$, where $\widehat{T}$ is either the age of death or the last year of the working stage of life-cycle (64 years old).[16] We use $r = 2\%$ for all calculations in this section and report the corresponding results when $r = 0\%$ in Appendix E.2.

	All	η_1	η_2	η_3
Over life cycle (21-death)				
Percentage of time in bad health	15%	58%	23%	4%
Income losses + total medical costs	\$1,511	\$8,896	\$1,935	\$225
(Percentage of average earnings) [a]	*(3.9%)*	*(23%)*	*(5%)*	*(0.6%)*
Over working age (21-64)				
Percentage of time in bad health	10%	55%	14%	1%
Income losses + total medical costs	\$1,031	\$7,147	\$1,201	\$76
(Percentage of average earnings) [a]	*(2.7%)*	*(18%)*	*(3%)*	*(0.2%)*

[a] Average earnings in our baseline model is \$38,648 per year.

Table 11: Annual monetary losses due to bad health (poor+fair). The top panel is over life cycle till death while the bottom panel is between 21 and 64 (working ages). The interest rate for computing the present value is 2%.

Table 11 displays the average lifetime costs of bad health over the entire life-cycle (top panel) and over the working period (bottom panel), starting from age 21. Two main points

[16] For the monetary losses over working age (21-64), $\widehat{T}$ is set to the age of death if an individual dies before reaching the age of 64 years old.

are worth noticing. First, on average, because of bad health realizations, people experience monetary losses over their entire life cycle of about $1,500 per year, and average losses over the working period of about $1,000 per year.

Second, the inequality in monetary losses across different health types is large. While people with the best health type (η_3) experience losses of only about $200 per year over the entire life-cycle, this number is close to $9,000 for people with the worst health type (η_1). The reason for this remarkable difference can be seen in the second row of the table: while people with the best health type spend only 4% of their lifetime being unhealthy, people with the worst health type find themselves in bad health for more than half of their life (58%). Thus, the large lifetime losses of the worst health type are driven by the long sickness spells which make up for a large portion of their relatively short lifespans.

	Over life-cycle (21-death)				Over working periods (21-64)			
	All	η_1	η_2	η_3	All	η_1	η_2	η_3
Annual monetary losses	$1,511	$8,896	$1,935	$225	$1,031	$7,147	$1,201	$76
Composition (%)								
- Medical costs paid by insurance	36%	33%	39%	39%	32%	33%	33%	18%
- Out-of-pocket medical costs	27%	22%	30%	36%	20%	20%	21%	11%
- Income losses	37%	45%	31%	24%	48%	47%	46%	71%

Table 12: Composition of annual monetary loss due to bad health (poor+fair) using 2% interest rate when computing the present value.

Next, we turn to decomposing the sources of the monetary losses due to bad health. Table 12 displays the distribution of these losses by three components: medical costs paid by insurance, out-of-pocket medical costs, and income losses. The left-hand-side panel in this table refers to the entire life-cycle, while the right-hand-side panel refers to the working stage only (age 21 to 64). In both cases, a substantial portion of the losses is due to income drops: it is the largest component of the losses that are not covered by insurance. In case of the working-age population, income losses represent almost half of total monetary losses due to bad health. Over the entire life-cycle the contribution of income losses due to bad health is about 40%. This difference is mainly due to the fact that medical spending increases quickly with age and thus plays a more important role for older people. Another important observation is that only about a third of monetary losses due to bad health are covered by insurance.

Overall, two important conclusions can be drawn from the results presented in this section. First, the lifetime monetary costs of bad health are substantial and those born with the worst health type have significantly higher costs over their life-cycle. Second, studies that confine the effects of bad health only to medical expenses significantly underestimate the total losses that unhealthy people experience over their lives.

5.3 The welfare losses due to bad health

Because bad health also affects one's disutility from working and life expectancy, welfare is a more comprehensive measure than the monetary costs of bad health. The first and last column of Table 13 suggest that the effects of bad health on welfare could be large because of its effects on life expectancy: people who do not experience any bad health shocks live, on average, six additional years. Hence, we now turn to computing how bad health affects one's welfare and through which channels.

	Baseline					When everyone is always healthy [a]
	all	η_1	η_2	η_3	variation due to η	
Average age at death	77.4	63.0	73.8	81.5	21 %	83.4

Table 13: Average age at death in the baseline case. In the counterfactual scenario, since everyone is healthy, the distribution of age at death is the same for all η_i.

To construct our welfare measure, we follow the same logic used when measuring the monetary costs: we compare each individual with his hypothetical self in a situation when his health is always good but everything else is the same. We then compute the percentage reduction in annual consumption that makes the person's welfare in the counterfactual environment to be the same as in the baseline.

Formally, the realized lifetime utility of an individual in the baseline is

$$U^{BS} = \sum_{t=1}^{\widehat{T}_d+1} \beta^t \Big(u(c_t^*, l_t^*, h_t) \times \mathbf{1}_{alive_t} + (1 - \mathbf{1}_{alive_t})\, \theta_{Beq} \frac{(k_t^* + k_{Beq})^{1-\rho}}{1-\rho} \Big),$$

where $\{c_t^*, l_t^*, k_t^*\}_{t=20}^{\widehat{T}_d}$ are optimal decisions, and $\widehat{T}_d$ is age at death. $\mathbf{1}_{alive_t}$ is an indicator function that equal one if a person is alive in period t and zero otherwise. We can similarly define the counterfactual lifetime utility of the same person in case of good health during each period.

The (realized) welfare costs of bad health is the reduction in consumption (λ_c) such that one's lifetime utility in the counterfactual is the same as that in the baseline:

$$U^{BS} = \sum_{t=1}^{\widehat{T}_d^H+1} \beta^t \Big(u\big((1-\lambda_c)\, c_t^{**}, l_t^{**}, h_t = G \big) \times \mathbf{1}_{alive_t} + (1 - \mathbf{1}_{alive_t})\, \theta_{Beq} \frac{(k_t^{**} + k_{Beq})^{1-\rho}}{1-\rho} \Big),$$

where $\{c_t^{**}, l^{**}, k_t^{**}\}$ and $\widehat{T}_d^H$ are the optimal decisions and age at death in case of unexpected continuous good health, respectively. To convert this magnitude to dollar values, we also report it as an annual consumption reduction or $\lambda_c \overline{c}^{**}$, where $\overline{c}^{**}$ is one's average consumption over the lifetime when drawing only good health realizations.

	all	η_1	η_2	η_3	β_L	β_H
Compensated consumption equivalence	\$1,933	\$6,380	\$2,690	\$854	\$1,866	\$2,018
(% consumption equivalence, λ_c)	*(10.6%)*	*(36.8%)*	*(14.8%)*	*(4.4%)*	*(10.3%)*	*(11%)*
Contribution (%)						
- Only medical expenses losses	25%	39%	22%	17%	24%	26%
- Only income losses	38%	57%	42%	9%	55%	17%
- Only non-monetary losses	44%	32%	33%	77%	14%	79%

Table 14: Welfare losses due to bad health (poor+fair). The dollar value is calculated from $\lambda_c \overline{c}^{**}$ where λ_c and $\overline{c}^{**}$ are the percentage of consumption reduction and average life time consumption of each individual when always healthy.

Table 14 displays the compensating consumption equivalent (CCE) averaged over all individuals and by health type and discount factor type.[17] The first line shows that, on average, CCE represents 10.6% of annual consumption (or \$1,933), thus indicating substantial welfare consequences of experiencing bad health. While the welfare losses vary little by discount factor type, there are remarkable differences by health types. Removing bad health realizations for people with the worst health type (η_1) is worth around a third of their annual consumption, while for people with the best health type (η_3) it is worth less than 5%. This is perhaps not surprising given the average person with the worst heath type spend more than half of their life in bad health.

The other lines of Table 14 report the contribution of various factors to the welfare losses of bad health. We compute them by performing three counterfactual experiments. In the first experiment, bad health *only* affects one's medical spending. In the second one, bad health *only* affects one's productivity and disutility from work, i.e., there is no effect on life expectancy or medical spending. In the third one, bad health *only* affects one's survival probability, i.e., individuals who become sick experience a decline in their life expectancy but no change in their productivity, disutility from work, or medical spending. In each of these three experiments, we recompute the corresponding compensating consumption equivalent by comparing one of the counterfactual baselines with the hypothetical situation of no bad health realizations. We report the resulting CCE as the percentage of the CCE corresponding to the situation when health affects people through all channels (first row of the table).[18]

The largest source of welfare losses due to bad health is its effect on life expectancy (44%), followed by its effect on income (38%). There is, however, substantial heterogeneity

[17] The dollar values of monetary and welfare losses are not directly comparable. To compute monetary losses, we discount the reduction in resources at each age by $1/(1+r) = 0.98$. To compute welfare losses, we discount the sum of lifetime utilities by discount factors 0.877 and 0.992 for inpatient and patient people, respectively.

[18] Our decomposition exercise is not supposed to sum to 100% by construction. The purpose of this exercise is to rank the importance of each channel through which health affects individuals.

in terms of the importance of different channels for people with different health types and discount factors.

While for people with the worst health type, almost 60% of welfare losses come from the income channel, for people with the best health type this number is only 9%. At the same time, while the survival channel contributes only a third to the welfare losses of people with the worst health type (η_1), it represents almost 80% of the welfare losses of those with the best health type (η_3). This difference is largely due to the correlation between health type and discount factor. In fact, the survival channel has little impact on the welfare of impatient people (only 14%), but plays a dominant role for patient people (almost 80%), as reported in the last two columns of this table. The latter is due to their different discount factor and hence valuation of events that happen later in life, like a reduction in life expectancy.

Another important dimension to consider is the concentration of losses due to bad health and the contribution of health types to its total variation. The first three columns of Table 15 display the percentage of aggregate losses experienced by people in the top 5%, 10%, and 20% of the distribution for monetary and welfare losses, respectively. For both measures, the distribution of lifetime costs of bad health is highly concentrated: for example, people at the top 10% of the distribution account for 56% and 42% of the monetary and welfare losses due to bad health, respectively. There is, however, a significant difference depending on the measure of costs used. Monetary costs are noticeably more concentrated: while the top 5% accounts for 24% of welfare losses, the top 5% accounts for almost 40% of monetary losses. The difference is much smaller for the top 20% of the distribution, thus indicating that this gap is largest among people at the very top of the losses distribution. The lower concentration of welfare losses is due to the fact that individuals' consumption is shielded from the highest monetary losses by means-tested insurance (guaranteed consumption floor).

	Over life cycle (21-death)			variation due to η
	top 5%	top 10%	top 20%	
Monetary losses (21-death)				
- Income losses + total medical costs	38%	56 %	75%	69%
Welfare losses				
- Compensated consumption equivalence	24%	42%	71%	30%

Table 15: Concentration of losses due to bad health (poor+fair) and variation of losses due to health types. The reported numbers in column 2 to 4 are in percentage of aggregate loss at top 5%, 10%, and 20%. For monetary loss, we use 2% interest rate when computing the present value and include the costs paid by insurance. The results when the costs paid by insurance are excluded are very similar.

The last column of Table 15 displays the contribution of health types η to the variance of the lifetime losses due to bad health. It shows that health types play an important role

in explaining the variance of lifetime losses: its contribution counts for 69% of the variation in monetary losses and for 30% in that of welfare losses.

The difference in the contribution of health types to the variance of two different measures of losses is due to the following. Health types account for 73% of the variation of the fraction of lifetime spent being unhealthy (see also second row of Table 11), this, in turn, is an important factor generating monetary losses. In contrast, the contribution of health types to life expectancy is smaller (about 20%, see Table 13). Because the largest contributor to welfare is life expectancy, health types play less of a role in explaining variation in welfare.

In sum, bad health generates large welfare costs, of the order of 10% of annual consumption on average. This happens because bad health lowers life expectancy and because of limited insurance opportunities against monetary losses. In addition, and importantly, the variation in welfare losses and (even more) in monetary losses, arises due to fixed health types, that is factors that are likely pre-determined earlier in life.

6 Conclusions and directions for future research

We develop a structural framework for understanding the differences in economic outcomes by health and for measuring the lifetime consequences of being unhealthy over the entire life-cycle. Our modeling approach emphasizes the complex nature of health dynamics and the role played by multidimensional innate differences across individuals.

We provide several important findings. Regarding our estimated health shock process, we find that health types are important drivers of health dynamics and that the variation in health transitions due to health types is much larger than that due to history-dependence. Regarding the implications of our estimated life-cycle model, we show that taking into account the correlated structure of innate differences across individuals is important to understand the disparity in economic outcomes by health. More specifically, the worse economic outcomes of the unhealthy are to a large extent due to the fact that, compared to the healthy, they are more likely to have lower fixed labor productivity and to be less patient.

We use our estimated model to evaluate the lifetime consequences of being unhealthy and we find these consequences to be substantial, both in monetary and welfare terms. On average, people lose about $1,500 per year over the life-course because of bad health realizations. In welfare terms, bad health realizations are, on average, equivalent to a 10% reduction in annual consumption. Our decomposition analysis shows that an important component of the monetary costs is the loss in labor income, especially for people of working age, while the effects of health on life expectancy are very important drivers of welfare costs. We also document that both measures of lifetime costs are very concentrated and unequally

distributed across health types and that the contribution of health types to the variation in the lifetime costs of bad costs is large, especially for monetary costs.

Our work indicates several directions for future research. A first one is to develop a structural framework that incorporates all mechanisms that can generate the interaction of health and economic outcomes. We allow for health to affect economic outcomes and for innate factors to affect both economic outcomes and health. The next step is to incorporate the possible effects of economic outcomes on health. This is a challenging extension since it raises several identification issues, which include determining what factors have to be included as inputs in health production function, and how long it takes for health investments to fully affect future health. The complex nature of the health dynamics that we document also emphasizes the role of fixed ex-ante heterogeneity. This should also be taken into account when modeling the health production technology.

A second important direction for future research is to better understand the direct effects of health on labor market productivity. Our results show that an important part of the monetary costs of bad health comes from its effect on labor income. It is possible that this occurs because health affects human capital accumulation. Modeling this relationship also raises several identification issues. For instance, how does health affect human capital? Is it through human capital investment or learning by doing? How should we capture the correlation between health types and fixed factors affecting labor productivity? In our approach, we estimate fixed labor productivity directly from the data. With endogenous human capital the issue of estimating ex-ante fixed factors affecting human capital accumulation and estimating their correlation with health types becomes very complex, especially since they must be estimated inside the structural model.

Finally, it is also important to understand the micro-foundations of the innate difference across individuals. Our estimates show that fixed characteristics, namely, health types, patience, and fixed labor productivity, are correlated. What are the mechanisms generating this correlation? What is the role of genetics and personality traits versus circumstances early in life? Understanding and modeling these issues requires incorporating the childhood stage of the life-cycle and modeling children's human capital formation through parental time and monetary investments. This approach also requires incorporating recent insights from personality psychology into a structural framework (see for instance Almlund et al., 2011, for a review). While challenging, the importance of this line of research for policy analysis cannot be underestimated: our results emphasize the importance of fixed innate characteristics for generating large losses due to bad health. Hence, developing policies to prevent these losses requires understanding the formation of these innate characteristics.

References

[1] Almlund, M., Duckworth, A. L., Heckman, J., Kautz, T. (2011). Personality psychology and economics. In E. A. Hanushek, S. Machin, L. Woessmann (Eds.), Handbook of the economics of education (Vol. 4., pp. 1-181). Elsevier

[2] Ameriks, J., Briggs, J., Caplin, A., Shapiro, M., Tonetti, C., 2020. Long-Term Care Utility and Late in Life Saving. Journal of Political Economy

[3] Anda, R., Felitti, V., Bremner, D., Walker, J., Whitfield, C., Perry, B., Dube, S., Giles, W., 2006. The enduring effects of abuse and related adverse experiences in childhood. A convergence of evidence from neurobiology and epidemiology. Eur Arch Psychiatry Clin Neurosci. 256(3), pp 174–186.

[4] Arthur, B., 1981. The Economics of Risks to Life. American Economic Review, 71.

[5] Attanasio, O., Kitao, S., Violante G., 2011. Financing Medicare: A General Equilibrium Analysis. In Shoven, J.B. (Ed.), Demography and the Economy, University of Chicago Press..

[6] Barth, D., Papageorge, N., Thom, K. 2020. Genetic Endowments and Wealth Inequality. Journal of Political Economy.

[7] Blundell, R., Borella, M., Commault, J., 2020a. "Why Does Consumption Fluctuates at Old Age and How Should the Government Insure It?" NBER working paper N 27348

[8] Blundell, R., Britton, J., Costa Dias, M., French, E., 2020b. "The Impact of Health on Labor Supply Near Retirement". Journal of Human Resources

[9] Bolt, U., 2021. What is the Source of the Health Gradient? The Case of Obesity. Mimeo, UCL

[10] Capatina, E., 2015. Life-cycle Effects of Health Risk. Journal of Monetary Economics, 74, pp.67-88.

[11] Capatina, E., Keane, M., Maruyama, S., 2021. Health and Earning Dynamics over the Life-Cycle. Mimeo, University of New South Wales.

[12] Carroll, C., 1997. Buffer Stock Saving and the Life Cycle/Permanent Income Hypothesis. Quarterly Journal of Economics CXII(1), pp 1–56.

[13] Carroll, C., Slacalek, J., Tokuoka, K., White, M., 2017. The Distribution of Wealth and the Marginal Propensity to Consume, Quantitative Economics, 8(3), pp.977-1020.

[14] Case, A., Lubotsky, D., Paxson, C., 2002. Economic Status and Health in Childhood: The Origins of the Gradient. American Economic Review, 92, pp.1308-1664.

[15] Case, A., Fertig, A., Paxson, C., 2005. The Lasting Impact of Childhood Health and Circumstance. Journal of Health Economics, 24(2), pp.365-389.

[16] Case, A., Paxson, C., 2010. Causes and Consequences of Early-Life Health. Demography, 47, pp.S65-S85.

[17] Cole, H., Kim, S., Krueger, D., 2018. Analyzing the Effects of Insuring Health Risks: On the Trade-off between Short-Run Insurance Benefits vs. Long-Run Incentive Costs. Review of Economic Studies (forthcoming).

[18] Conesa, J., Costa, D., Kamali, P., Kehoe, T., Nygard, V., Raveendranathan, G., Saxena, A., 2018. Macroeconomic Effects of Medicare, Journal of the Economics of Ageing, 11, pp.27-40.

[19] Conti, G., Heckman, G., 2010. Understanding the Early Origins of the Education–Health Gradient: A Framework That Can Also Be Applied to Analyze Gene–Environment Interactions. Perspectives on Psychological Science, 5(5), pp 585–605.

[20] Cronqvist, H., Siegel, S., 2015. The Origins of Savings Behavior, Journal of Political Economy, 123(1), pp.123-169.

[21] Currie, J., and Madrian., B., 1999. "Health, Health Insurance, and the Labor Market". In Handbook of Labor Economics, Volume 3, Ed. By Ashenfelter O., and Card, D.

[22] Cutler, D., Lleras-Muney, A., and Vogl, T., 2011. Socioeconomic Status and Health: Dimensions and Mechanisms. In Glied, S. and Smith, P. (Eds), Oxford Handbook of Health Economics, Oxford University Press.

[23] De Nardi, M., 2004. "Wealth Inequality and Intergenerational Links". Review of Economic Studies, 71, pp.743-768.

[24] De Nardi, M., French, E., Jones, J., 2010. Why Do the Elderly Save? Journal of Political Economy, 118(1), pp.39-75.

[25] De Nardi, M., French, E., Jones, J., 2016. Medicaid Insurance in Old Age. American Economic Review, 106(11), pp.3480-3520

[26] Dobkin, C., Finkelstein, A., Kluender, R., Notowidigdo, M., 2016. The Economic Consequences of Hospital Admissions, NBER Working Papers.

[27] Erosa, A., Fuster, L., Kambourov, G., 2016. Towards a Micro-founded Theory of Aggregate Labour Supply. Review of Economic Studies, 83(6), pp.1001-1039.

[28] Epper, T., Fehr, E., Fehr-Duda, H., Kreiner, C., Lassen, D., Leth-Petersen, S., Rasmussen, G., 2018. Time Discounting and Wealth Inequality, Working paper.

[29] Fang, H., Krueger, D., 2021. The Affordable Care Act after a Decade: Its Impact on the Labor Market and the Macroeconomy.

[30] Felitti, V., Anda, R.,Nordenberg, D., Williamson, D., Spitz, A., Edwards, V., Koss, M., Marks, J., 1998. Show less Relationship of Childhood Abuse and Household Dysfunction to Many of the Leading Causes of Death in Adults. The Adverse Childhood Experiences (ACE) Study. American Journal of Preventive Medicine, 14(4)

[31] Finkelstein, A., Luttmer, E., Notowidigdo, M., 2009. Approaches to Estimating the Health State Dependence of the Utility Function, American Economic Review: Papers & Proceedings, 99(2), pp.116-121.

[32] Floden, M., 2008. A Note on the Accuracy of Markov-chain Approximations to Highly Persistent AR(1) Processes. Economic Letters, 99(3), pp.516-520.

[33] French, E., 2005. The Effects of Health, Wealth, and Wages on Labor Supply and Retirement Behavior. Review of Economic Studies, 72(2), pp.395-427.

[34] French, E., Jones, J., 2011. The Effects of Health Insurance and Self-Insurance on Retirement Behavior. Econometrica, 79(3), pp.693-732.

[35] Gouveia, M. and Strauss, R. P., 1994. Effective Federal Individual Tax Functions: An Exploratory Empirical Analysis, National Tax Journal, 47(2), pp.317-39.

[36] Guvenen, F., 2009. An empirical investigation of labor income processes, Review of Economic Dynamic, 12(1), pp.58-79.

[37] Hai, R., Heckman, J., 2021. A Dynamic Model of Health, Addiction, Education, and Wealth. Mimeo, University of Miami

[38] Hall, R., Jones, C., 2007. The Value of Life and the Rise in Health Spending. Quarterly Journal of Economics, 122(1), pp.39-72.

[39] Halliday, T., 2008. Heterogeneity, State Dependence and Health. Econometrics Journal, 11(3), pp.499-516.

[40] Harris et al., 2017. Molecular genetic contributions to self-rated health

[41] Hendricks, L., 2007a. Retirement Wealth and Lifetime Earnings, International Economic Review, 48(2), pp.421-456.

[42] Hendricks, L., 2007b. How Important is Discount Rate Heterogeneity for Wealth Inequality?, Journal of Economic Dynamics & Control 31, pp.3042-3068.

[43] Hosseini, R., Kopecky, K., Zhao, K., 2021a. The Evolution of Health over the Life Cycle. Review of Economic Dynamics.

[44] Hosseini, R., Kopecky, K., Zhao, K., 2021b. How Important is Health Inequality for Lifetime Earnings Inequality? Mimeo, Federal Reserve Bank of Atlanta

[45] Huggett, M., Ventura, G., Yaron, A., 2010. Sources of Lifetime Inequality. American Economic Review, 101(7), pp.2923-2954.

[46] Idler, E., and Benyamini, Y., 1997. Self-Rated Health and Mortality: A Review of Twenty-Seven Community Studies. Journal of Health and Social Behavior, Vol 38, pp.21-37.

[47] Jeske, K., Kitao, S., 2009. U.S. Tax Policy and Health Insurance Demand: Can a Regressive Policy Improve Welfare? Journal of Monetary Economics, 56(2), pp.210-221.

[48] Jung, J., Tran, C., 2016. Market Inefficiency, Insurance Mandate and Welfare: U.S Health Care Reform 2010. Review of Economic Dynamics, Vol 20, 132-159.

[49] Kaiser Family Foundation, 2004., The Cost of Care for the Uninsured. What Do We Spend, Who Pays, and What Would Full Coverage Add to Medical Spending?. Available at http://www.kff.org/uninsured/upload/The-Cost-of-Care-for-the-Uninsured-What-Do-We-Spend-Who-Pays-and-What-Would-Full-Coverage-Add-to-Medical-Spending.pdf.

[50] Kaplan, G., 2012. Inequality and the Life Cycle. Quantitative Economics, 3, pp.471–525.

[51] Keane, M., Wolpin, K., 1997. The Career Decisions of Young Men. Journal of Polictical Economy, 105(3), pp.473-522.

[52] Krueger, D., Mitman, K., Perri, F., 2016. Macroeconomics and Household Heterogeneity. In Taylor, J., Uhlig, H., (Eds.), Handbook of Macroeconomics Volume 2, Elsevier.

[53] Krusell, P., Smith, A., 1998. Income and Wealth Heterogeneity in the Macroeconomy, Journal of Political Economy, 106(5), pp.867-896.

[54] Lange, F., McKee, D., 2012. The Evolution of Latent Health over the Life Course. Mimeo, Yale University.

[55] Lawrance, E., 1991. Poverty and the Rate of Time Preference: Evidence from Panel Data, Journal of Political Economy, 99, pp.54-77.

[56] Lockwood, L., 2018. Incidental Bequests and the Choice to Self-Insure Late Life Risks. American Economic Review, 108(9), pp 2513-2550.

[57] Lundborg, P., Nilsson, M., Vikstrom, J., 2015. Heterogeneity in the Impact of Health Shocks on Labour Outcomes: evidence from Swedish workers. Oxford Economic Papers, 67(3), pp.715-739.

[58] Murphy, K., Topel, R., 2006. The Value of health and Longevity, Journal of Political Economy, 114(5), pp.871-904.

[59] Nakajima, M., Telyukova, I., 2020. Home Equity in Retirement. International Economic Review.

[60] Ozkan, S., 2014. Preventive vs. Curative Medicine: A Macroeconomic Analysis of Health Care over the Life Cycle. Mimeo, University of Toronto.

[61] Papageorge, N., Thom, K., 2019. Genes, Education, and Labor Market Outcomes : Evidence from the Health and Retirement Study. Journal of the European Economic Association.

[62] Pashchenko, S., 2013. Accounting for Non-annuitization. Journal of Public Economics 98, pp.53-67.

[63] Pashchenko, S., Porapakkarm, P., 2013. Quantitative Analysis of Health Insurance Reform: Separating Regulation from Redistribution. Review of Economic Dynamics 16, pp.383-404.

[64] Pashchenko, S., Porapakkarm, P., 2016a. Cross-subsidization in Employer-based Health Insurance and the Effects of Tax Subsidy Reform. National Tax Journal, 69(3), pp.583-612

[65] Pashchenko, S., Porapakkarm, P., 2016b. Medical Spending in the U.S.: Facts from the Medical Expenditure Panel Survey Dataset. Fiscal Studies, 37(3-4), pp.689-716.

[66] Pashchenko, S., Porapakkarm, P., 2017. Work Incentives of Medicaid Beneficiaries and the Role of Asset Testing. International Economic Review, 58(4), pp.1117-1154.

[67] Pijoan-Mas, J. and Ríos Rull, V., 2014. Heterogeneity in Expected Longevities, Demography, 51(6), pp.2075-2102.

[68] Parro, F., and Pohl, V., 2017. Health Shocks, Human Capital, and Labor Market Outcomes. Mimeo, University of Georgia.

[69] Poterba, J., Venti, S., Wise, D., 2017. The Asset Cost of Poor Health, The Journal of the Economics of Ageing, 9, pp.172-184.

[70] Quadrini, V., 1999. The importance of entrepreneurship for wealth concentration and mobility. Review of Income and Wealth, 45(1), pp 1–19.

[71] Robinson, L., 2007. How US Government Agencies Value Mortality Risk Reductions, Review of Environmental Economics and Policy, 1(2), pp.283-299.

[72] Romeis,J., Scherrer, J., Xian, H., Eisen, S., Bucholz, K., Heath, A., Goldberg, J., Lyons, M., Henderson, W., True, W., 2000. Heritability of Self-Reported Health. Health Services Research, Dec;35(5 Pt 1), pp 995-1010.

[73] Rust, J., Phelan, C, 1997. How Social Security and Medicare Affect Retirement Behavior In a World of Incomplete Markets, Econometrica, 65(4), pp.781-831.

[74] Samwick, A., 1998. Discount Rate Heterogeneity and Social Security Reform, Journal of Development Economics, 57(1), pp.117-146.

[75] Shepard, D, Zeckhauser, R., 1984. Survival versus Consumption. Management Science, Vol. 30, No 4.

[76] Silventoinen K, Posthuma D, Lahelma E, Rose RJ, Kaprio J., 2007. Genetic and environmental factors affecting self-rated health from age 16-25: a longitudinal study of Finnish twins. Behavioral Genetics, 37, pp 326-33

[77] Smith, J., 1999. Healthy Bodies and Thick Wallets: The Dual Relation Between Health and Economic Status. Journal of Economic Perspectives, 13(2), pp.144-166.

[78] Storesletten, K., Telmer, C., Yaron, Y., 2004. Consumption and Risk Sharing Over the Life Cycle. Journal of Monetary Economics, 51(3), pp.609-633.

[79] Van Doorsaler, E., and Gerdtham, U., 2002. Does Inequality in Self-Assessed Health Predict Inequality in Survival by Income? - Evidence from Swedish Data. Social Science and Medicine, 57, pp.1621-1629.

[80] Viscusi, K., 1993. The Value of Risks to Life and Health. Journal of Economic Literature, 31(4), pp.1912-1946.

[81] Warner, J., Pleeter, S., 2001. The Personal Discount Rate: Evidence from Military Downsizing Programs, American Economic Review, 91, pp.33-53.

Online appendix

A The data

We use three data sets: the Panel Study of Income Dynamics (PSID), the Health and Retirement Study (HRS), and the Medical Expenditure Panel Survey (MEPS). The PSID tracks individuals over a long period of time and contains excellent information, for instance, on self-reported health status, and on labor supply during the entire working stage. The HRS has a large sample size and a lot of information on people age 50+, including wealth, health, and labor supply. The MEPS contains high-quality information about both total and out-of-pocket medical spending. Thus, using all three data sets allows us to exploit the advantages of all of them and to construct the best possible data for the United States.

For each data set, we select a sample of male household heads with 12 to 14 years of education (that is, with a high school degree or at most 2 years of college). We normalize all nominal variables to the 2013 base year using the Consumer Price Index (CPI).

The PSID is a nationally representative panel that surveys individuals and their families. It started in 1968, on an annual basis, but has been administered bi-annually since 1997. Individuals' self-reported health is available from 1984. We use all available waves from 1984 to 2017. To construct our panel data sample, we include individuals who do not have missing observations on self-reported health status and assign the longitudinal weight in 2013 (our base year) as individual weight.[19] This gives us a sample of 2,038 individuals or 26,194 individual-wave observations (on average, individuals are observed for 12.9 waves). To construct the bi-annual panel data for the maximum likelihood estimation of our two-year health shock process, we further drop every other wave prior to 1997.

In addition to using the PSID to estimate the health shock process that we describe in Section 2, we also use it for both our first and second steps estimation. At the first step, we use the PSID to estimate the annual labor income shock process (which we model as an AR(1)). We describe our estimation of the annual income process and its conversion to the bi-annual discretized process to be used in our model in Appendix D.3.

At the second step, we use the PSID to construct the targeted moments for the estimation of our life-cycle model listed in Section 4.1, namely, labor income, employment and wealth. The wealth data is available only in the 1994 and 1999-2017 waves.

We use the HRS to estimate health-dependent survival probabilities and survival after age 50. When possible, we also use the HRS to validate moments from the PSID. In addition, we use the HRS to validate our estimated model. We use 1994-2017 waves of the RAND

[19] Our results are robust to using equal weights.

HRS Longitudinal file 2016 (v1).

We use MEPS for our first step estimation. Specifically, we use it to estimate total medical expense shocks, coverage of employer-sponsored (ESHI) and individual health insurance, and the probability of obtaining ESHI. The medical spending reported in MEPS is cross-checked with insurers and providers and is thus very accurate.[20] We use waves 1999/2000-2016/2017 of MEPS.

B Additional details on our health process estimation

This section provides additional information regarding our paper's empirical findings. Specifically, it provides more details on the estimation of the health shock process described in Section 2 and explains how we construct our health transition probabilities and likelihood function. It also provides several estimation results, including our survival probabilities, health process estimation using different samples, and the full set of transition probabilities. Finally, it illustrates the implications of our health shock process in terms of the distribution of people by length of unhealthy spells.

B.1 Constructing the two-year health transition probabilities

The two-year transition probabilities in Figure 1 are constructed as follows. Denote h_{it} as health status of an individual i at age t. The probability of moving to good health conditional on currently being in bad health can be expressed as

$$\frac{\sum_i \mathbf{1}(h_{it}{=}B \ \cap \ h_{it+1}{=}G)}{\sum_i \mathbf{1}(h_{it}{=}B \ \cap \ h_{it+1}{=}\{B,G\})}.$$

The term $\mathbf{1}(\cdot)$ is the indicator function, which is equal to one if its argument is true and zero otherwise.

To construct the health transition profiles in Figure 2, denote the sequence of health statuses of an individual i in the past τ periods up to age t as h_{it}^{τ}. For people age 30 to 54 years old, we compute the probability of moving to good health during the next two years $(t+1)$, conditional on being unhealthy for at least τ consecutive periods as follows:

$$\frac{\sum_{t=30}^{54} \sum_i \mathbf{1}\left(h_{it}^{\tau}{=}B \ \cap \ h_{it+1}{=}G\right)}{\sum_{t=30}^{54} \sum_i \mathbf{1}\left(h_{it}^{\tau}{=}B \ \cap \ h_{it+1}{=}\{B,G\}\right)}.$$

[20] Pashchenko and Porapakkarm (2016b) provide more details on the MEPS dataset.

B.2 Likelihood function for the two-year health shock process

Consider an individual i whose health is observed for consecutive J periods. Denote as t_0 the earliest age of an uninterrupted health sequence till age $t_0 + J$: $\{h_{i,t_0}, h_{i,t_0+1}, ...h_{i,t_0+J}\}$. If an individual is of η_j-type, the conditional probability of the observed health sequence can be constructed from Eqs (2)-(5) and the estimated conditional survival probability in Eq. (7). Denote the conditional probability of this health sequence as $Pr\left(h_{t_0+\tau}.., h_{t_0+J} \mid h_{i,t_0}..., h_{i,t_0+\tau-1}, \eta_j\right)$.

Since one's health type η_j is unobserved, we compute one's expected likelihood function using

$$\sum_{j=1}^{3} Pr\left(\eta_j|\gamma_i, \mathbf{X}_{i,t_0}^{\eta}\right) \times Pr\left(h_{i,t_0+\tau}, .., h_{i,t_0+J} \mid h_{i,t_0}..., h_{i,t_0+\tau-1}, \eta_j\right),$$

where $Pr\left(\eta_j|\gamma_i, \mathbf{X}_{i,t_0}^{\eta}\right)$ is derived from the cumulative probability in Eq. (6). The probability above independent across individuals. Hence, we can write the overall log-likelihood as the sum of individuals' log-likelihoods,

$$L\left(\Theta\right) = \sum_{i=1}^{N} \log\Big(\sum_{j=1}^{3} Pr\left(\eta_j|\gamma_i, \mathbf{X}_{i,t_0}^{\eta}\right) \times Pr\left(h_{i,t_0+\tau}, .., h_{i,t_0+J} \mid h_{i,t_0}..., h_{i,t_0+\tau-1}, \eta_j\right)\Big),$$

where Θ is the set of parameters in Eqs (2)-(6).

B.3 Estimating the two-year survival probabilities

Figure B1 shows our estimated two-year survival probabilities from the HRS. Several observations are in order. First, one's survival probability decreases in health and the gap between people in poor and fair health is significantly larger than that between people in fair and good health. Second, one's health history does not matter much for people currently in poor or fair health, while it does matter for those whose current health is good. That is, people older than 70 who spend at least three periods in good health have noticeably higher chances to survive than their counterparts who have recently recovered from bad health and have thus been in good health for just one or two periods.

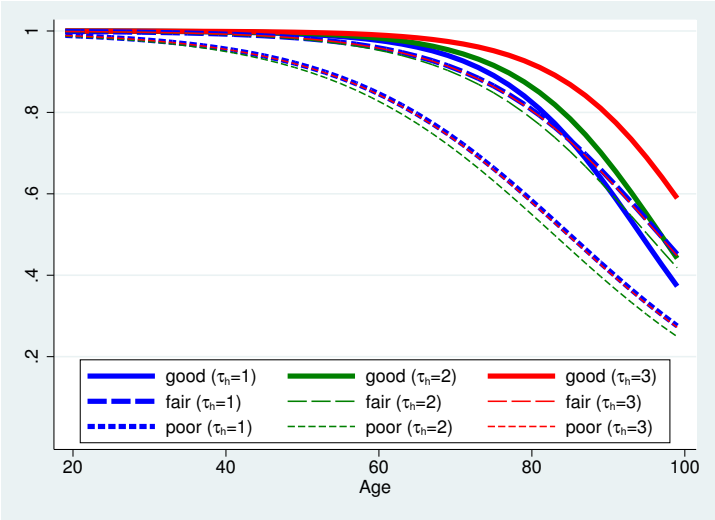


Figure B1: Estimated two-year survival probabilities by health status and health history.

B.4 Estimating the two-year health shock process (restricted sample)

Tables B1 and B2 report the estimation results from Eqs. (2)-(6) when we restrict the sample to individuals whose age t_0 is less than or equal to 39. The average t_0 in the restricted sample is between 25 and 30 years old, depending on the specification of T. These tables shows that the results for this more restricted sample are similar to those that we use for our main specification.

	T=5	T=4	T=3	T=2	T=1
Coefficients of history-dependence terms and health types in Eq. (2)-(3)					
a_2^B	0.015	0.030	0.136	0.441**	–
a_3^B	0.565	0.612*	0.543**		
a_4^B	0.045	0.179			
a_5^B	0.289				
$a_{\eta_1}^B$	2.347***	2.143***	1.915***	1.698***	1.886***
$a_{\eta_3}^B$	-2.901***	-2.870***	-2.36***	-2.388***	-1.799***
Coefficients of history-dependence terms and health types in Eq. (4)-(5)					
a_2^G	-0.487*	-0.283	-0.436*	-0.795***	–
a_3^G	-0.558*	-0.433	-1.215***		
a_4^G	-1.401***	-1.687***			
a_5^G	-1.979***				
$a_{\eta_1}^G$	4.532***	4.418***	4.090***	3.683***	1.715***
$a_{\eta_3}^G$	-1.417***	-1.567***	-2.304***	-2.741***	-2.771***
N	6064	7212	8668	9982	11346

Table B1: Estimation of the health shock process when the sample includes only individuals whose age t_0 is less or equal to 39 years old. $a_{\eta_2}^B$ and $a_{\eta_2}^G$ are normalized to zero. Being in bad/good health for one period ($\tau_B = 1, \tau_G = 1$) is set as the base case. (* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$)

	T=5	T=4	T=3	T=2	T=1
age t_0	-0.162*	-0.089*	0.038	-0.005	-0.026
$h_{t_0} = G$	-0.795	-0.995*	-1.696***	-1.892***	-2.182***
$h_{t_0} = P$	1.466	1.507	1.936	0.866	-0.376
2^{nd} tercile of γ	-0.250	-0.357	-0.557*	-0.494*	-0.679***
3^{rd} tercile of γ	-1.716***	-1.865***	-1.788***	-1.529***	-1.546***
2^{nd} quintile of k_{t_0}	-0.277	-0.0637	-0.0546	-0.830**	-0.669**
3^{rd} quintile of k_{t_0}	-0.380	-0.233	-0.300	-0.348	-0.578**
4^{th} quintile of k_{t_0}	-1.289	-0.853	-0.667*	-0.753**	-0.677**
5^{th} quintile of k_{t_0}	-2.192**	-1.345**	-1.566***	-1.547***	-1.431***

Table B2: Estimation results of the equation predicting health type in Eq. (6) when the sample includes only individuals whose age t_0 is less or equal to 39 years old. Fair health status, 1^{st} tercile of γ and 1^{st} quintile of k_{t_0} are set as the base case for the corresponding dummy variables. All estimations include dummy variables for 10-years bracket of birth year. (* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$)

.

B.5 Estimating the two-year health shock process (all education groups)

Education is a common determinant of many important outcomes. To further investigate the relationship between health types and education, we expand our PSID and HRS samples to include all education groups, namely, less than high school, high school, and at least college. We then repeat the estimation steps in Section 2.3 and include education dummy variables as an additional covariate in Eq (6). We report the results in Table B3 and B4.

The results confirm our finding that education is insufficient to capture the variation in health types that we document. Moreover, these additional estimates confirm our estimated relationship between fixed productivity and wealth quintiles, in the context of a larger sample size.

Similar to our estimation when using only high school group, Table B4 shows that people with a higher fixed productivity and higher initial wealth are less likely to be of the worst health type. The last two rows of Table B4 also highlight that having higher education is a significant predictor of better health types.

	T=5	T=4	T=3	T=2	T=1
Coefficients of history-dependence terms and health types in Eq. (2)-(3)					
a_2^B	0.554***	0.495***	0.507***	0.821***	–
a_3^B	1.181***	1.094***	1.206***		
a_4^B	0.933***	1.236***			
a_5^B	1.562***				
$a_{\eta_1}^B$	1.011***	1.204***	1.099***	1.324***	1.909***
$a_{\eta_3}^B$	-0.418	-0.646	-0.464	-0.732*	-1.419***
Coefficients of history-dependence terms and health types in Eq. (4)-(5)					
a_2^G	-0.432**	-0.363**	-0.355**	-0.729***	–
a_3^G	-0.499**	-0.382*	-1.140***		
a_4^G	-1.039***	-1.452***			
a_5^G	-1.856***				
$a_{\eta_1}^G$	1.404***	1.669***	1.772***	1.914***	1.967***
$a_{\eta_3}^G$	-1.292***	-1.775***	-1.966***	-2.486***	-2.964***
N	17811	19261	21918	23745	25593

Table B3: Estimation of the health shock process when the sample includes all education groups. $a_{\eta_2}^B$ and $a_{\eta_2}^G$ are normalized to zero. Being in bad/good health for one period ($\tau_B = 1, \tau_G = 1$) is set as the base case. (* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$)

	T=5	T=4	T=3	T=2	T=1
age t_0	-0.140*	-0.081**	-0.018	-0.013	-0.022
$h_{t_0} = G$	-1.952***	-2.150***	-2.469***	-2.598***	-2.452***
$h_{t_0} = P$	0.099	0.275	3.059	0.425	-0.042
2^{nd} tercile of γ	-1.418**	-1.024***	-1.045***	-0.949***	-0.869***
3^{rd} tercile of γ	-2.457**	-1.750***	-1.805***	-1.721***	-1.741***
2^{nd} quintile of k_{t_0}	-0.608	-0.304	-0.172	-0.274	-0.278
3^{rd} quintile of k_{t_0}	-1.016*	-0.429	-0.465	-0.660**	-0.713***
4^{th} quintile of k_{t_0}	-1.614**	-0.919**	-0.794**	-0.846***	-0.830***
5^{th} quintile of k_{t_0}	-2.955**	-1.421***	-0.925**	-0.899***	-0.729***
high school group	-1.362*	-0.881*	-0.734**	-0.568*	-0.450*
college group	-3.369**	-2.295***	-1.966***	-1.641***	-1.414***

Table B4: Estimation results of the equation predicting health type in Eq. (6) when the sample includes all education groups (less than high school, high school, and college). Less than high school, fair health status, 1^{st} tercile of γ and 1^{st} quintile of k_{t_0} are set as the base case for the corresponding dummy variables. All estimations include dummy variables for 10-years bracket of birth year. (* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$)

.

B.6 The two-year health transition probability ($T = 3$)

Table B5 reports the initial joint distribution of health status and fixed productivity taken from our PSID sample for people between the ages of 19 and 24 and shows that worse initial health and low fixed productivity are positively correlated.

	γ_L	γ_M	γ_H
$Pr\left(h_{21} = F, \tau_{21} = 1 \mid \gamma\right)$	0.084	0.062	0.026
$Pr\left(h_{21} = G, \tau_{21} = 3 \mid \gamma\right)$	0.926	0.938	0.973

Table B5: Joint distribution between health status and fixed productivity (h_{21}, τ_{21}) between age 19 and 24. The term τ_{21} represents the number of consecutive periods that an individual has been in health status h_{21}. The terms $\{\gamma_L, \gamma_M, \gamma_H\}$ refer to the three fixed productivity terciles.

Figure B2 reports our estimated transition probabilities for all health states when T=3. It features very large differences by health types. Figure B3 shows that our health transition probabilities, together with the initial measure of health types (η) in Table 3, reproduce the fractions of people in fair and poor health by age that we observe in the data.

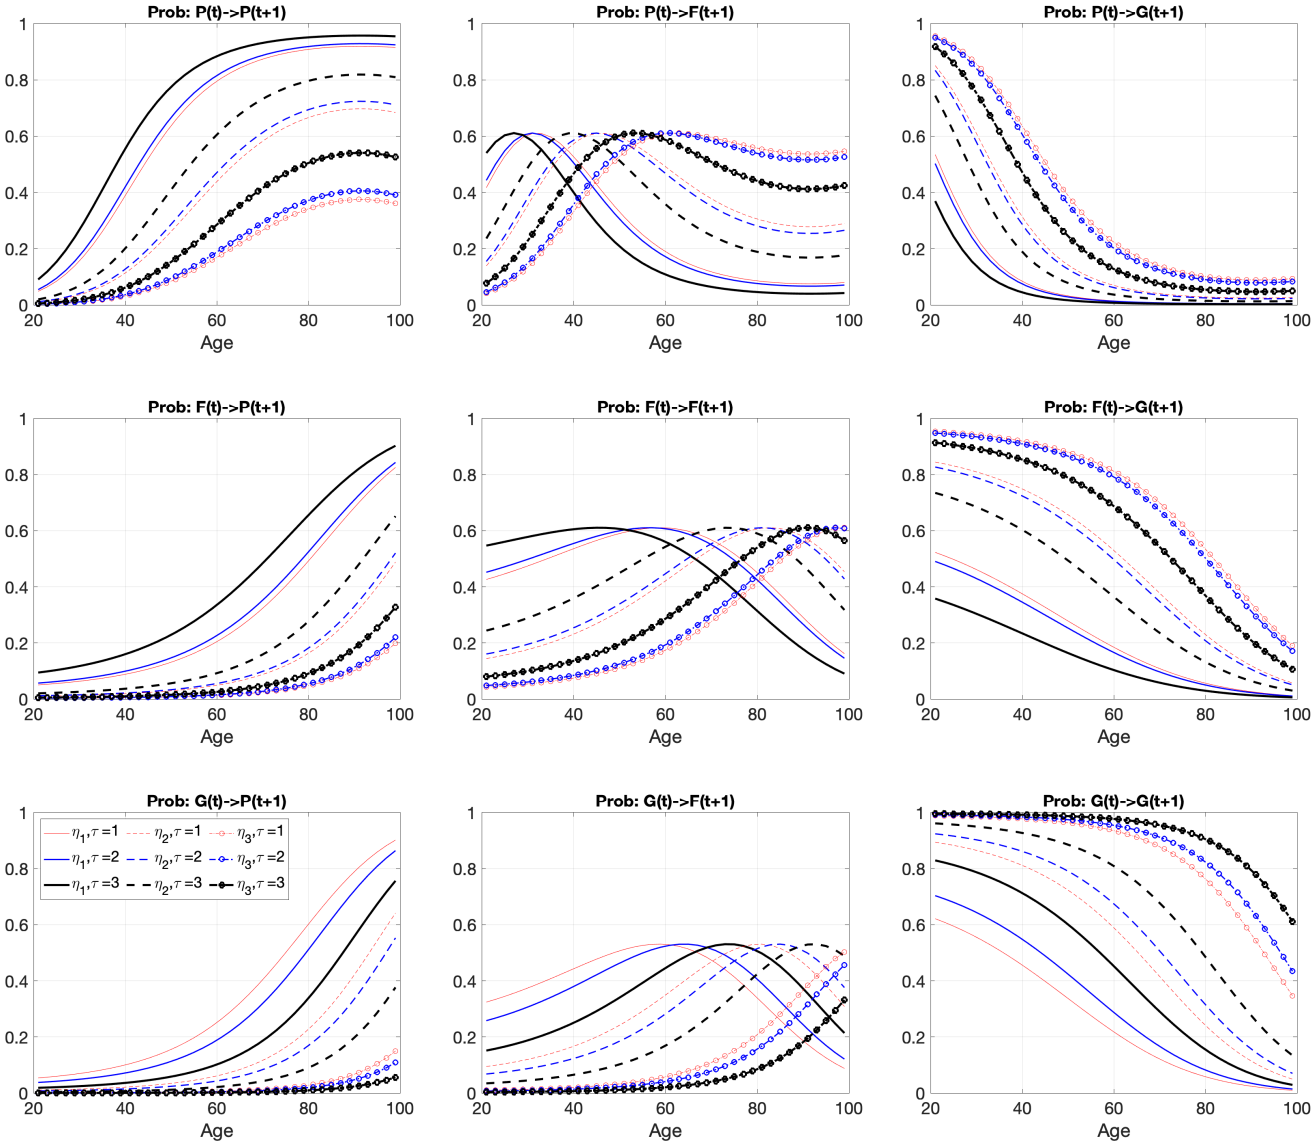


Figure B2: Estimated two-year transition probabilities ($T = 3$). Different line styles represent three health types: solid line for η_1-type, dashed line and line with dots for η_2- and η_3-type, respectively. Black, blue, and red lines correspond to the number of consecutive periods in the current health status (τ). The panels in the first, second, and third rows correspond to the case when individuals are currently in poor, fair, and good health, respectively. The first, second, and third columns correspond to poor, fair, and good health in the next two years.

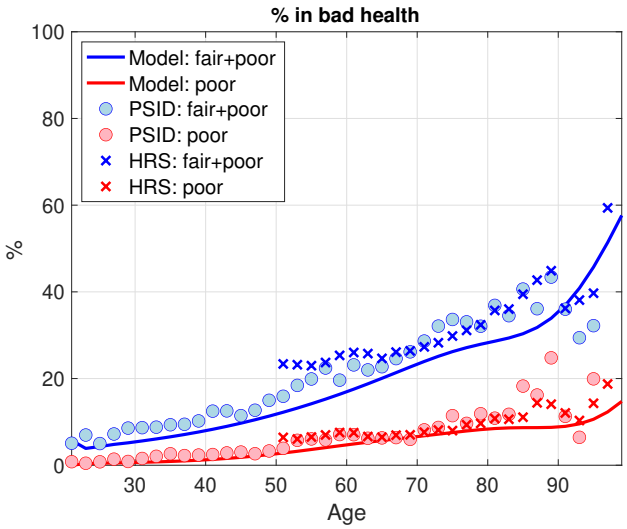


Figure B3: Fraction of people in bad health

B.7 What accounts for the long spells of bad health?

We also use our estimated bi-annual model of health dynamics to compute the distribution of the number of unhealthy periods over one's lifetime. The left graph of Figure B4 plots the distribution of people by the total number of periods spent unhealthy between the age of 21 and 80, while the dashed line in the right panel shows the corresponding cumulative distribution. Most people are relatively healthy over their life course: 70% of people experience no more than 5 periods of bad health. However, a non-trivial number of people spend more than a third of time being unhealthy between the age of 21 and 80: almost 10% of people experience 10 or more periods in bad health (note that the number of periods between the age of 21 and 80 is 30).

The right graph of Figure (B4) illustrate how this distribution differs across health types by comparing two groups: the solid line with circles refers to people of the worst health type (η_1), while the dotted one refers to those with better health types (η_2 and η_3). Among η_2- or η_3-individuals (solid line), almost no one experiences more than 15 unhealthy periods between the age of 21 and 80. In contrast, slightly more than 30% of η_1-people are unhealthy for 15 periods or longer. Thus, even though the measure of η_1-people is small (8.3% at age 21), they primarily account for the long right tail of the unhealthy period distribution in the left panel. In other words, long spells of bad health are mostly due to fixed heterogeneity, rather than to repeated bad realizations from a persistent health shock.

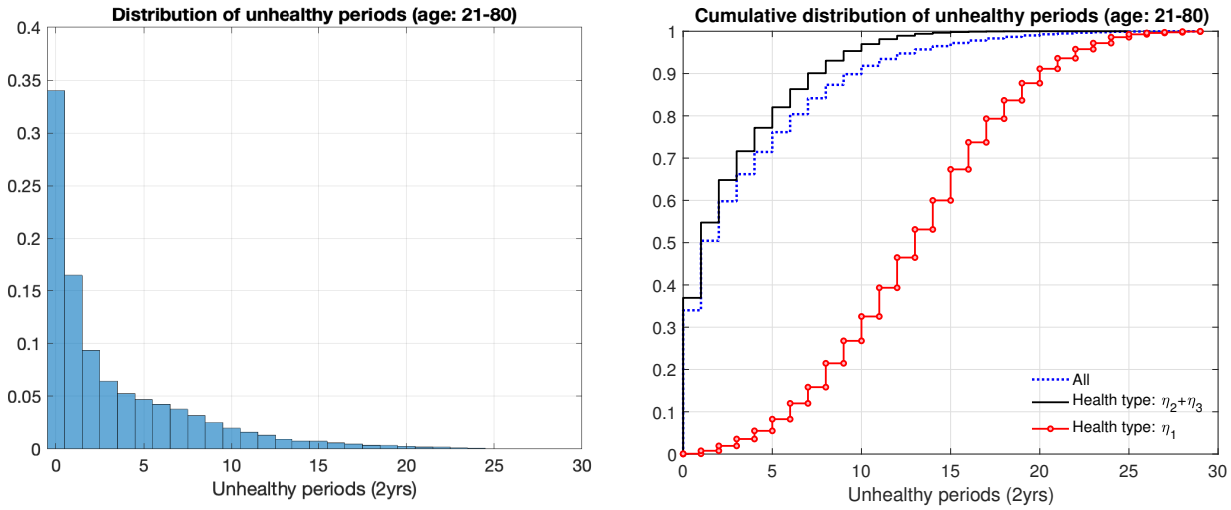


Figure B4: Distribution by lifetime unhealthy periods (age 21-80). Left panel: distribution among all individuals. Right panel: cumulative distribution of individuals with different heath types.

C Health types' interpretation

C.1 Polygenic scores and health outcomes in the HRS

In Section 2.5, we report the relationship between health and individuals' fixed characteristics using a balanced panel of individuals observed consecutively between age 55/56 and 65/66. In Table 4, the sample size for individuals with 0-1 periods being unhealthy is between 794 and 904, depending on the variable, while the sample size for 2-3 periods and 4-5 periods are 123-140 and 54-63, respectively. For the polygenic scores in Table 5, the sample size for individuals with 0-1, 2-3, and 4-5 periods being unhealthy are 611, 73, and 31, respectively.

In this section, we further investigate the correlation between health and genetic endowments. To do so, we use the same HRS sample as in Table 5 in the main text, and report the $25^{th}, 50^{th}$, and 75^{th} percentiles of four polygenic scores among individuals with different number of unhealthy periods. The resulting pattern is similar to the one reported in Table 5, where we only reported average polygenic scores. Among people reporting 4-5 unhealthy periods, all percentiles of the polygenic score for educational attainment are noticeably lower. In contrast, all percentiles of the polygenic scores associated with unhealthy behaviors (smoking and BMI) are higher for this group (one exception is the 25^{th} percentile of polygenic score for BMI). All percentiles of polygenic score predicting longevity among those reporting 4-5 unhealthy periods are also consistently lower.

# unhealthy periods	Polygenic scores (HRS)			
	educational attainment	smoking	BMI	longevity
	25^{th} *percentile*			
0-1	-0.761	-0.692	-0.665	-0.761
4-5	-1.610	-0.487	-0.725	-0.950
	50^{th} *percentile*			
0-1	-0.120	0.036	0.003	-0.028
4-5	-0.682	0.127	0.215	-0.332
	75^{th} *percentile*			
0-1	0.473	0.720	0.666	0.560
4-5	-0.075	0.734	1.089	0.486

Table C6: The $25^{th}, 50^{th}$, and 75^{th} percentiles of various polygenic scores by the number of unhealthy periods between ages 57 and 66. All individuals are healthy at age 55-56.

C.2 Empirical studies supporting the existence of health types

In this section we review the empirical literature studying the role of pre-determined factors in adult health. These studies can be broadly divided into two groups.

The first group of studies focuses on genetic contributions to health. Romeis et al. (2000) use several thousands of male-male twin pairs from the Vietnam Era Twin Registry. They find that the genetic contribution to adult health is 40%. Silventoinen et al. (2007) study a sample of Finnish twins and conclude that heritability accounts for 33% of variation in self-reported health at the age of 25. Studies that use actual genetic data find smaller but still significant contribution of genetic factors to health. For example, Harris et al. (2017) show that 13% of the variation in self-reported health can be explained by all common genetic variants.

The second group of studies focuses on the contribution of early childhood circumstances to adult health. Case et al. (2005) show that each chronic condition at age 7 raises the probability of reporting a chronic condition at age 42 by 4%. Moreover, if the condition is still present at age 16, the effect is twice as large. Conti et al. (2010) show that childhood health has a significant effect on the probability of having poor health at age 30, even controlling for cognitive and non-cognitive ability. Campbell et al. (2014) use biomedical data and show that early childhood intervention can significantly improve objective health measures in adulthood. A number of studies investigate the role of Adverse Childhood Experiences (ACE), such as having experienced or witnessed physical or mental abuse. They show that ACE has long-lasting effects: individuals with high ACE score when children have substantially worse health in their middle and old age (Anda et al., 2006; Felitti et al., 1998).

Taken together, this (growing) evidence shows that genetic factors and early childhood circumstances have significant impact on adult health, giving additional support to our findings that fixed ex-ante heterogeneity is an important determinant of health dynamics.

D First step estimation details

D.1 Medical shocks and insurance coverage

To estimate medical expenses, we follow Pashchenko and Porapkkarm (2017). That is, we first convert medical expenses in the MEPS to 2013 price using the CPI. Second, we compute total medical expenses for each individual over two-year periods. Third, we separate our sample into 12 age groups (20-24, 25-29, 30-34, ..., 75+), where we assign the age of each group to the mid-point of the corresponding age interval. For example, 22 for 20-24, 27 for 25-29, 32 for 30-34, etc. Then, for every age group, conditional on health status in the first year (poor, fair, good), we divide the two-year medical expenses into 3 bins: the bottom 50%, 50-90%, and the top 90%. After computing average of medical expenses in each bin, we multiply them by 1.60 for people younger than 65 years old and by 1.90 for people 65 or older to make medical spending in our model consistent with the aggregate medical spending in the National Health Expenditure Accounts (NHEA). We then fit the resulting adjusted medical expenses with a quadratic function of age. Figure (D5) shows the medical costs for each grid separately for each health status.

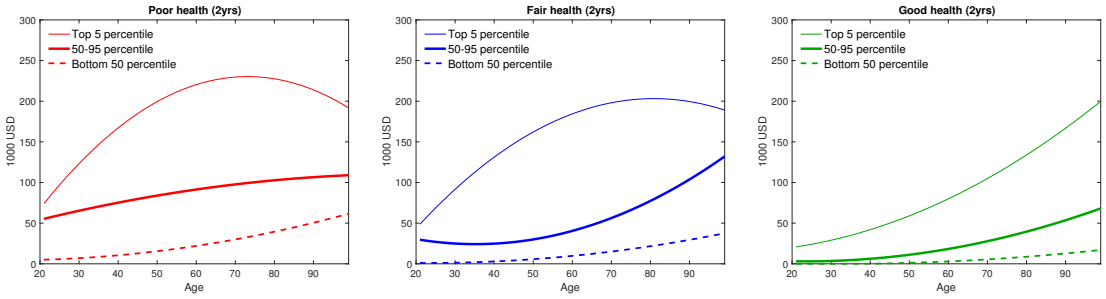


Figure D5: Two-year Medical expense grids by health status, x_t^h.

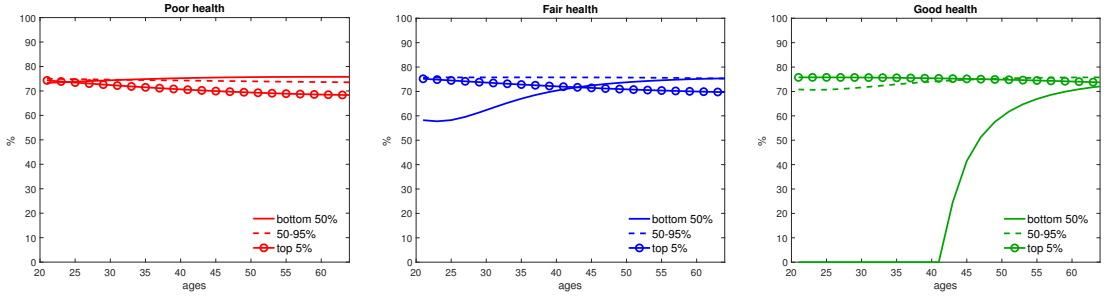


Figure D6: Private health insurance coverages: $cvg\left(x_t^h, i_H\right), i_H \in \{1, 2\}$.

To determine the fraction of medical expenses covered by private insurance $cvg(x_t^h, i_H)$ where $i_H \in \{1, 2\}$, we do the following. We estimate medical expenses paid by private

insurers as a function of total medical expenses and year dummies using only individuals who are categorized as individually insured or group-insured. Then, we convert our estimates into the fraction of expenses covered by insurers. Figure D6 shows the estimated coverage by medical expense grids.

For the parameters related to private health insurance market, we use Pashchenko and Porapakkarm (2017)'s estimates and set the proportional load ξ to 1.07, and the fixed loads φ^h to \$100 for the healthy and \$2,100 for those in fair or poor health (annually).[21]

D.2 Tax system

We set the Medicare, Social Security, and consumption tax rates to 2.9 percent, 12.4 percent, and 5.67 percent, respectively. We use the Social Security rules for 2013, hence we set the maximum (annual) taxable income for Social Security ($\overline{y}_{ss}$) to \$113,700. We set the capital tax rate to 36%, as Holter et al. (2019).

For the progressive labor income tax function, we follow Holter et al. (2019) whose estimates of $a_{\tau 0}$ and $a_{\tau 1}$ depend on family structure. Since the average family size $\overline{n}_t$ in our model ranges from 1.9 to 3.2, we set $a_{\tau 0}$ and $a_{\tau 1}$ to 0.940772 and 0.158466, respectively, which corresponds to their estimates for married families with one child.[22]

D.3 Stochastic labor productivity

We use the PSID (annual data before 1997 and bi-annual after 1997) to estimate our labor income shock process. We define workers as described in Section 2.2 and compute their annual labor income, defined as earnings plus income from business. We have 17,277 individual-wave observations from 1,730 individuals whom we observe working at least in two waves. We obtain earning residuals ($\gamma_i + u_{it}$) from a fixed effect regression of Eq. (1).[23]

Next, we construct the empirical autocovariance matrix of the earning residuals and estimate the parameters of the productivity shock by minimizing the distance between the

[21] Since the loads existing in individual health insurance market are unobservable, we use the indirect estimates obtained by Pashchenko and Porapakkarm (2017). In their model with rich representation of the US health insurance market these loads are identified from the observed purchase of individual health insurance.

[22] Since income in our model is over two years, we convert it into annual income before applying the tax function. Then we convert the resulting tax obligations into two-year payments. We use the same approach for the Social Security benefits.

[23] Note that the parameters in Equation (13) are assumed to be independent of health status and age. Since most workers are healthy and over 90% of healthy people work, we are less concerned about the selection problem when estimating the parameters of Υ_t directly from the data. An alternative approach is to use only the sample of healthy workers younger than 60, but this would reduce our sample size.

empirical autocovariance matrix and the corresponding matrix implied by Eq.(13).[24] Our resulting estimates of annual labor income shock are $\rho_\nu = 0.947$, $\sigma^2_\varepsilon = 0.02$, $\sigma^2_{\nu_0} = 0.09$, and $\sigma^2_\gamma = 0.051$, and are within the range of values estimated in the literature. We then use the estimated annual AR(1) process to simulate annual income shock for a large number of individuals. From this simulated data, we construct age-dependent transition matrix for two-year labor income shock with 9 grid points, equally-spaced and expanding with age. We discretize the fixed productivity into three terciles $\{\gamma_L, \gamma_M, \gamma_H\}$.

D.4 Summary of the parametrization of the baseline model

Table D7 below summarizes parameters of our life-cycle model.

Parameter name	Notation	Value	Source
Parameters set outside the model			
Risk aversion	ρ	3.0	
Average family size	$\overline{n}_t$	1.9-3.2	PSID
Tax function parameters	$a_{\tau 0}$	0.940772	Holter, et al. (2019)
	$a_{\tau 1}$	0.158466	"
capital income tax	τ_k	36%	"
consumption tax	τ_c	5%	"
Medicare premium (per year)	P_{MCR}	\$1,055	
Labor productivity			
- Persistence parameter (annual)	ρ_ν	0.9472	PSID
- Variance of innovations (annual)	σ^2_ε	0.0198	"
- Initial distribution	$\sigma^2_{\nu_0}$	0.093	"
- Fixed effects	σ^2_γ	0.051	"
Proportional load in ind ins	ξ	1.07	
Fixed load in ind ins (per year)	φ^h		
- healthy		\$100	Pashchenko and Porapakkarm (2017)
- unhealthy (poor+fair)		\$2,100	"
Parameters used to match some targets			
Discount factors (per year)	$\beta_{low}, \beta_{high}$	0.877,0.992	wealth profiles
% individual with β_{low} by η at 20	$Pr(\beta_{low}\|\eta_1)$	77.8%	"
	$Pr(\beta_{low}\|\eta_2)$	79.4%	"
	$Pr(\beta_{low}\|\eta_3)$	38.0%	"
Bequest parameters			
- Strength	θ_{Beq}	1,905	"
- Shifter	k_{Beq}	\$182,707	"
Consumption floor (per year)	$\overline{c}$	\$3,505	"
Per-period utility of being alive	$\overline{b}$	7.149	VSL of \$2M

Table D7: Parameters of the baseline model

[24] This is a standard procedure commonly used in the literature. See for example, Storesletten et al. (2004) and French (2005).

E Additional results

In this section, we provide supplementary discussions and additional results about the implications of our estimated structural model. More specifically, we discuss our results about the estimated heterogeneity in discount factors and compare them with estimates from other studies. We also report the monetary costs of bad health computed using an alternative interest rate and the welfare costs of bad health when the VSL is higher than in our baseline estimation.

E.1 Discussion of the estimated heterogeneity in the rate of time preferences.

Our estimated discount factors are 0.877 and 0.992 for impatient and patient groups, respectively. In this section, we compare the difference in the rates of our time preferences with those from other studies that allow for patience heterogeneity and structurally estimating them to match wealth moments.

Among the structural studies that allow for heterogeneity in discount rates, we can distinguish two groups. The first group, starting from the seminal paper of Krusell and Smith (1998) and including Hubmer et al. (2019) and Krueger et al. (2016), shows that a small difference in discount rates is enough to generate "enough" wealth inequality. These studies use as relevant moments Gini coefficient or wealth holdings among different percentiles of the wealth distribution and feature households that are infinitely lived (or age stochastically and thus can also potentially live infinitely).

The second group of studies find that significantly larger difference in patience are needed to match the data. Crawford and O'Dea (2020) estimate a structural model on linked survey and administrative data in the UK and find substantial heterogeneity in discount rates: the bottom/top 10% of the estimated discount rate distribution is equal to 0.98/1.125, respectively. French and Jones (2011) use a rich model to account for retirement and saving decision after age 50 (estimated using the HRS data) and estimate discount factors varying from 0.8 to more than 1 for different groups of people. Hendricks (2007), which, as the first group of papers focuses on the impact of discount factor heterogeneity on wealth inequality, estimates discount factor varying across people from 0.91 to more than 1 in his benchmark model. Moser and Olea de Souza e Silva (2019) use the HRS and the CPS find discount rates varying from 0.905 to 0.999 between the 10th and the 90th percentile of their distribution. These studies target wealth moments *conditional on age* in their estimation. A common modeling feature of these studies is a deterministic (non-stochastic) life-cycle framework and the presence of a bequest motive.

Because we model a deterministic life-cycle structure and bequest motives, and we target wealth inequality evolution over the life-cycle, our paper belongs to the second group of studies.

The intuition why in our kind of framework the estimated heterogeneity in discount rates tends to be larger is due to the absence of impatient households according to Carroll (1997)'s definition among certain age groups. Let us elaborate a bit on this point.

Based on Carroll (1997), in a standard model without uncertainty and with no income growth (income is the same and equal to y in the current and future periods), households are impatient if the following is true:

$$u'(y) \geq \beta(1+r)u'(y),$$

which corresponds to the Euler equation of an individual who chooses to set his savings to 0. This leads to the impatience condition $\beta(1+r) \leq 1$. In our framework with bequest motive and survival uncertainty, this impatience condition can be represented as follows:

$$u'(y) \geq \beta(1+r)(\zeta u'(y) + (1-\zeta)v'(0)),$$

where $v'(0)$ is the marginal utility of leaving no bequests and ζ is survival probability. Using our parametrization of bequest and utility functions, the expression above can be written as follows

$$1 \geq \beta(1+r)\left[\zeta + (1-\zeta)\ \eta\left(\frac{\phi}{y}\right)^{-\sigma}\right].$$

Consider, for example, a relatively young individual with survival probability ζ equal to 0.99. In the absence of a bequest motive, an individual would be considered impatient in our model if his discount factor is below 0.99 (since our interest rate is 2%). Using our estimates of the bequest motive, the threshold β that makes such an individual impatient becomes 0.83, assuming that he has average income. As one's survival probability decreases and income increases with age, the cutoff that defines impatience goes down. This means that after a certain age, all individuals become patient.

When all individuals in a certain group are at the same side of the impatience cutoff, it is harder to generate wealth inequality compared to the situation in which some individuals are patient and some are impatient. We are facing this situation since we are targeting wealth inequality by age. Thus, we need a larger heterogeneity in discount factors to generate difference in saving behavior among individuals of a certain age since they are all patient under the modified definition of patience, that is the one that accounts for bequest motives.

In terms of interpretation, in our framework β represents certain characteristics that are fixed ex-ante and affect saving behavior. They can be interpreted as people's non-cognitive abilities and, more precisely, as the ability to delay gratification. Research in personality psychology starting from the seminal work of Mischel et al. (1989) show that the patience or ability to delay gratification measured in childhood is significantly correlated with outcomes later in life.

While the economic mechanism linking patience or ability to delay gratification and saving behavior is well-understood, it is less clear how exactly this type of non-cognitive abilities affects health and labor market outcomes, once controlling for education. However, there is evidence that such a relationship does exist. For example, Golsteyn et al. (2014) using a survey data in Sweden linked to administrative records show that there exists substantial adverse relationship between high rate of time preferences measured at age 13 and health and labor outcomes later in life.

E.2 Monetary loss of bad health when interest rate is low

To compute the monetary costs of bad health we use the following formula:

$$loss_i = \frac{1}{\widehat{T}} \sum_{t=1}^{\widehat{T}} \frac{y_{it}^H - y_{it}^{BS}}{(1+r)^t}.$$

where y_{it}^{BS} and y_{it}^H are income net of total medical spending of an individual i at time t in the baseline and counterfactual case with no bad health realizations, respectively. In the main text, we use the interest rate of 2% to compute the losses. In this section, we recompute the losses using the interest rate of 0%. The results of this exercise are reported in Tables E8 and E9. The overall losses are larger when lower interest rate is used, but the decomposition exercise reveals the same pattern regarding the importance of income and medical spending losses. The corresponding concentration of monetary losses and its variation due to health type are reported in the upper part of Table E11. The overall patterns are similar to the case when using 2% interest rate.

E.3 Alternative calibration with VSL of $6 millions

In our baseline parameterization, we adjust the scaling parameter $\overline{b}$ so that the value of statistical life (VSL) among working age population implied by our model is $2 million dollars. In this section, we report the results from an alternative parameterization when $\overline{b}$ is set to match a VSL of $6 millions. Note that all other parameters in the model are the

	All	η_1	η_2	η_3
Over life cycle (21-death) using 0% interest rate				
% of time in bad health	15%	58%	23%	4%
Income losses + total medical costs	\$3,303	\$16,800	\$4,613	\$627
(% of average earnings) [a]	*(8.5%)*	*(43%)*	*(12%)*	*(1.6%)*
Over working age (21-64) using 0% interest rate				
% of time in bad health	10%	55%	14%	1%
Income losses + total medical costs	\$2,026	\$13,267	\$2,569	\$130
(% of average earnings) [a]	*(5.2%)*	*(34%)*	*(6.6%)*	*(0.3%)*

[a] Average earnings in our baseline model is \$38,648 per year.

Table E8: Annual monetary loss due to bad health (poor+fair) using 0% interest rate. The top panel is over life cycle till death while the bottom panel is between 21 and 64 (working ages).

	Over life-cycle (21-death)				Over working periods (21-64)			
	All	η_1	η_2	η_3	All	η_1	η_2	η_3
Annual monetary losses	\$3,303	\$16,800	\$4,613	\$627	\$2,026	\$13,267	\$2,569	\$130
Composition (%)								
- Medical costs paid by insurance	35%	32%	36%	43%	28%	30%	26%	19%
- Out-of-pocket medical costs	28%	21%	32%	41%	18%	18%	19%	10%
- Income losses	37%	47%	32%	15%	54%	51%	55%	70%

Table E9: Composition of annual monetary loss due to bad health (poor+fair) using 0% interest rate when computing the present value.

same as in the baseline. The welfare losses when the targeted VSL is set to a higher value are reported in Table E10. Not surprisingly, the importance of non-pecuniary consequence of bad health increases, compared with the baseline case. This is because with higher VSL, life is more valuable and there are larger welfare costs of bad health coming from a shorter lifespan.

The lower part of Table E11 reports the concentration of welfare losses and its variation due to health types. A comparison with Table 15 in the main text reveals that the increase in VSL does not affect the concentration of welfare losses. In contrast, a larger VSL noticeably lowers the contribution of health types toward the variance of welfare losses because the survival channel becomes a larger fraction of welfare losses and health types contribute relatively less to variation in life expectancy.

	all	η_1	η_2	η_3	β_L	β_H
Compensated consumption equivalence	\$2,320	\$6,416	\$3,025	\$1,321	\$1,860	\$2,896
(% consumption equivalence, λ_c)	*(12.6%)*	*(37.2%)*	*(16.6%)*	*(6.8%)*	*(10.3%)*	*(15.5%)*
Contribution (%)						
- Only medical expenses losses	22%	40%	20%	12%	24%	20%
- Only income losses	32%	58%	38%	6%	56%	13%
- Only non-monetary losses	60%	47%	47%	86%	20%	92%

Table E10: Welfare loss due to bad health (poor+fair) using VSL=\$6M. The dollar value is calculated from $\lambda_c \overline{c}^{**}$ where λ_c and $\overline{c}^{**}$ are the percentage of consumption reduction and average life time consumption of each individual when always being healthy.

	Over life cycle (21-death)			variation due to η
	top 5%	top 10%	top 20%	
Monetary losses (21-death) using 0% interest rate				
- Income losses + total medical costs	32%	49 %	71%	69%
Welfare losses using VSL=\$6M				
- Compensated consumption equivalence	23%	43%	74%	17%

Table E11: Concentration of losses due to bad health (poor+fair) and variation due to health types. The reported numbers in column 2 to 4 are in percentage of aggregate loss at top 5%, 10%, and 20%. For monetary loss, we use 0% interest rate when computing the present value and include the costs paid by insurance. (The results when excluding insurance are similar.) The welfare loss is based on the case when the VSL is calibrated to \$6M.