

The Lifetime Costs of Bad Health

Mariacristina De Nardi

University of Minnesota, Federal Reserve Bank of Minneapolis, CEPR, and NBER, USA

Svetlana Pashchenko

University of Georgia, USA

and

Ponpoje Porapakkarm

National Graduate Institute for Policy Studies, Japan

First version received August 2018; Editorial decision September 2023; Accepted August 2024 (Eds.)

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.

Key words: Health, Health insurance, Medical spending, Wealth–health gradient, Life-cycle models

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

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 (own calculations, see Section 5.1 for details).

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: (1) health changes economic circumstances; (2) economic circumstances affect health, and (3) healthy and unhealthy people

The editor in charge of this paper was Dirk Krueger.

are ex-ante different. While we adopt the first mechanism, the novelty of our approach is that we also allow for the third mechanism. In fact, modelling the latter is where we see a big gap in the literature. While the existence of the third mechanism is well-recognised, most structural models adopt a simple structure in which all ex-ante 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 modelling of ex-ante differences, which are likely multidimensional, and possibly correlated with each other.

Our focus on the role of factors that are pre-determined 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 (Felitti *et al.*, 1998; Case *et al.*, 2005; Anda *et al.*, 2006; Case and Paxson, 2010; Conti and Heckman, 2010; Cronqvist and Siegel, 2015; Harris *et al.*, 2017; Barth *et al.*, 2020). 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 (unhealthy). We formulate and estimate a parametric model of health shocks that allows for both history dependence and fixed heterogeneity, and that matches the patterns that we observe. Finally, we incorporate our estimated health process in a life-cycle model. The ex-ante differences across individuals in our model are three-dimensional and include fixed heterogeneity in health (or health types), in labour 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 emphasise 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 labour 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 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 process is correlated with fixed heterogeneity in labour 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 (in 2013 USD) 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 labour 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 realisations. 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, labour 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, 2022, for an extensive review) or the contribution of health to inequality in earnings (Capatina *et al.*, 2021; Hosseini *et al.*, 2021). The findings of these studies crucially depend on what generates difference in economic outcomes by health: a direct link between the two or ex-ante differences across individuals. We contribute to the literature by emphasising 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 emphasise the importance of three points: (1) how we measure health, (2) how best to describe health dynamics, and (3) 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.¹ In addition, several 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, Van Doorsaler and Gerdtham, 2002, and Pijoan-Mas and Ríos Rull, 2014 for a more recent examination). Finally, and very importantly, this measure is available for a long period of time in the Panel Study of Income Dynamics (PSID) (and the Health and Retirement Study (HRS)), which makes it ideally suited to investigate the lifetime costs of bad health.

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.* (2020), Hosseini *et al.* (2022), and 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. Importantly, Blundell *et al.* (2023) point out that properly accounting for ex-ante differences across individuals is a more important modelling issue than the choice of health measure. Our estimation provides new insights about the role of these ex-ante 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.*, 2020; Hosseini *et al.*, 2021). 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, labour productivity, and other economic variables (Jesse and Kitao, 2009; Pashchenko and Porapakkarm, 2013, 2017; Conesa *et al.*, 2018). The second approach is to assume that health is endogenous and can be affected by medical spending (Jung and Tran, 2016; Ozkan, 2023), effort (Cole *et al.*, 2019) or healthy/unhealthy habits (Bolt, 2021; Hai and Heckman, 2021).² In modelling the direct link between health and economic

1. 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.

2. 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 health transitions.

outcomes, we follow the first approach, but we emphasise 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 labour market outcomes (see [Currie and Madrian, 1999](#), for a review). Several structural studies investigate the importance of health for economic decisions such as savings ([Hubbard *et al.*, 1994](#); [De Nardi *et al.*, 2010, 2016](#); [Lockwood, 2018](#); [Ameriks *et al.*, 2020](#); [Nakajima and Telyukova, 2020](#)) or labour supply ([French, 2005](#); [French and Jones, 2011](#); [Capatina, 2015](#)). 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 pre-determined by factors formed at birth or early age is very important for analysing 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' maximisation of expected lifetime utility. Early examples include [Arthur \(1981\)](#), [Rosen \(1988\)](#), and [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 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 organised as follows. Section 2 documents empirical facts related to health dynamics and estimates the health processes. Section 4 introduces our life-cycle model and Section 5 describes its estimation. We present the results and conclusions in Section 6 and Section 7, respectively.

2. HEALTH AND LABOUR PRODUCTIVITY

For our health process estimation, we primarily use the PSID. When possible, we also use the HRS to cross-check moments from the PSID and to validate our estimated model. We report more detail about these datasets, our samples, and how we use them in [Appendix A, Supplementary Material](#).

For each dataset, we use a sample of male household heads with 12–14 years of education (corresponding to the high-school degree or at most 2 years of college education). We normalise 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 labour 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, e.g. [Rust and Phelan, 1997](#); [French, 2005](#); [Capatina, 2015](#)) 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 centre and right panels display the 2-year health transition probabilities. Details of their construction are in [Appendix B.1, Supplementary Material](#).





FIGURE 1
Health

Notes: Left panel: percentage of people in bad health by age. Middle panel: percentage of people moving from bad to good health in the next 2 years. Right panel: percentage of people moving from good to bad health in the next 2 years. “Model” refers to the simulated data from our estimated health process described in Section 2.3

(a)



(b)




FIGURE 2

Dynamics of health conditional on duration of previous health status. (a) Percentage of transitions from bad to good health conditional on being in bad health and (b) percentage of transitions from good to bad health conditional on being in good health

Notes: Each period is every 2 years. “Model” refers to the simulated data from our estimated health process described in Section 2.3

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.

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 2 years, conditional on being in bad (good) health for at least τ consecutive periods (with a period being 2 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.³

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; 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 labour productivity

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

To compute fixed labour productivity, we start by computing real labour income from the PSID for workers younger than 70, who work at least 520 h per year, and earn at least the federal minimum wage. As French (2005), we then estimate the age profile of labour income by health status among working individuals using the following fixed-effect regression:

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

where inc_{it} is persons' i labour income at age t , d_t^j is the coefficient on a dummy variable capturing the interaction between one's age and health status, and $\hat{\gamma}_i$ is unobserved fixed labour 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 .

3. This negative duration dependence is a robust pattern. It also holds when we use the annual data (using the 1984–97 waves of the PSID). See De Nardi *et al.* (2017) for those results. It is also present when we exclude people ever receiving Social Security Disability Insurance.



FIGURE 3

Fixed labour 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 labour 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 labour 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 labour 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 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 heterogeneity in labour productivity. A similar approach is taken by [Low and Pistaferri \(2015\)](#), who allow fixed productivity to affect the stochastic evolution of disability shocks.

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 three-state health 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 in during the current period. In terms of 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 to preserving tractability, combining the history of poor and fair health avoids 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)

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

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

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

where Λ 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 $f_{age}^P(t)$ and $f_{age}^F(t)$, 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: $a_{\eta_1}^B > a_{\eta_2}^B > a_{\eta_3}^B$. 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

$$Pr(P_{t+1} | G_t, \tau_G, \eta) = \Lambda \left(\sum_{\tau=1}^{T-1} a_{\tau}^G \mathbf{1}_{(\tau_G=\tau)} + a_T^G \mathbf{1}_{(\tau_G \geq T)} + f_{age}^G(t) + a_{\eta}^G D_{\eta} \right), \quad (4)$$

$$Pr(F_{t+1} \cup P_{t+1} | G_t, \tau_G, \eta) = \Lambda \left(\sum_{\tau=1}^{T-1} a_{\tau}^G \mathbf{1}_{(\tau_G=\tau)} + a_T^G \mathbf{1}_{(\tau_G \geq T)} + f_{age}^G(t) + b_2 + a_{\eta}^G D_{\eta} \right), \quad (5)$$

where the type-dependent coefficients a_{η}^G 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 $a_{\eta}^B = a_{\eta}^G = 0$ 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:

$$Pr(\eta \leq \eta_j | \mathbf{X}_{t_0}^{\eta}) = \begin{cases} \Lambda(b_{\eta_j} + \mathbf{B}_{\eta} \times \mathbf{X}_{t_0}^{\eta}) & \text{for } j = \{1, 2\}, \\ 1 & \text{for } j = 3, \end{cases} \quad (6)$$

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, $\mathbf{X}_{t_0}^{\eta}$ 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}_{η} are the corresponding coefficients (we provide more details on the construction of the log-likelihood function in [Appendix B.2, Supplementary Material](#)).

The variables that we include in $\mathbf{X}_{t_0}^{\eta}$ are age t_0 , birth-year dummies (in 10-year windows), health status h_{t_0} , fixed labour 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 window. 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 modelling survival. We do so by specifying the following logit model for the 2-year survival probability of an individual at age t , who has been in health state h_t for τ_h periods,

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

The first two terms in the logit function capture history dependence. Similar to the health 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, Supplementary Material](#) shows our estimated 2-year survival probabilities.

After estimating our survival probabilities, we estimate equations (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 equations (2) and (3) we impose $a_{\eta_1}^B > a_{\eta_2}^B > a_{\eta_3}^B$ where $a_{\eta_2}^B$ is normalised to zero. Similarly, in equations (4) and (5) $a_{\eta_2}^G$ is normalised 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 $\{a_{\eta_1}^G, a_{\eta_2}^G, a_{\eta_3}^G\}$. Finally, a cumulative probability must be monotonically increasing; thus, b_1 in equation (3) and b_2 in equation (5) must be non-negative, and $b_{\eta_1} \leq b_{\eta_2}$ in equation (6).

It is worth noting that the identification of fixed heterogeneity comes from the dependence of the health transition probabilities on how long one has been in the current health state. That is, in a health process in which we allow for T lags, in the absence of health types, the implied

TABLE 1
Estimation results for the health process in equations (2) and (3) in the top panel and equations (4) and (5) in the bottom panel

	$T = 5$	$T = 4$	$T = 3$	$T = 2$	$T = 1$
Coefficients of history-dependence terms and health types in equations (2) and (3)					
a_2^B	0.075	0.071	0.129	0.288	-
a_3^B	0.826***	0.810**	0.675**		
a_4^B	0.527	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 equations (4) and (5)					
a_2^G	-0.391*	-0.366*	-0.369*	-0.770***	-
a_3^G	-0.241	-0.183	-1.086***		
a_4^G	-1.007***	-1.691***			
a_5^G	-1.921***				
$a_{\eta_1}^G$	4.527***	3.786***	1.637***	1.806***	2.006***
$a_{\eta_3}^G$	-1.447***	-1.639***	-2.317***	-2.555***	-2.871***
N	9,028	9,765	11,126	12,096	13,083

Notes: The columns refer to specifications controlling for different number of lags of past health. The terms $a_{\eta_2}^B$ and $a_{\eta_2}^G$ are normalised to zero. Being in bad/good health for one period ($\tau_B = 1$, $\tau_G = 1$) is the base case. All estimations include a quadratic in age whose coefficients depend on current health status (poor, fair, good). * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

TABLE 2
Estimation results to predict health type in equation (6)

	$T = 5$	$T = 4$	$T = 3$	$T = 2$	$T = 1$
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***
$h_{t_0} = P$	1.462	2.072*	2.409	2.385	1.021
Second tercile of γ	-0.247	-0.337	-0.508**	-0.546**	-0.642***
Third tercile of γ	-1.203***	-1.374***	-1.189***	-1.286***	-1.355***
Second quintile of k_{t_0}	-0.001	-0.129	-0.048	-0.459*	-0.469*
Third quintile of k_{t_0}	-0.621	-0.429	-0.367	-0.378	-0.601**
Fourth quintile of k_{t_0}	-0.748	-0.606	-0.691*	-0.701**	-0.761***
Fifth quintile of k_{t_0}	-2.347***	-1.616***	-1.169***	-1.280***	-1.265***

Notes: The columns refer to specifications controlling for different number of lags of past health. Fair health status, first tercile of γ and first quintile of k_{t_0} are the base for the corresponding dummy variables. All estimations include dummy variables for 10-years windows of birth year. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

health transition probabilities are the same for people who spend T or more periods in particular health state. Thus, the lower probability of exiting the current health status after T lags in the data is attributed to the effect of fixed heterogeneity.

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 ($a_{\eta_1}^B$, $a_{\eta_3}^B$, $a_{\eta_1}^G$, $a_{\eta_3}^G$) are statistically significant.

Hence, even when we allow for five 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 artefact 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.

It is worth noticing that, for ease of interpretation, Tables 1 and 2 report our parameters' statistical significance without taking into account that one's productivity and survival probabilities are estimated in a previous step and thus subject to uncertainty. To quantify the effects of this simplification, Appendix B.4, [Supplementary Material](#) reports the results when we bootstrap to compute the 90% and 95% confidence intervals for our estimated health process and health types parameters.⁴ This check shows that the inference drawn in this case is consistent with the one in our baseline case.

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 equation (6). It shows that one's observed health status h_{t_0} , fixed labour 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 (second row), he is significantly less likely to be of η_1 -type. The opposite is true if his health is poor (third 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 labour 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 (*i.e.* age 39 and younger).⁵ These results are similar to those that we obtain from the full sample. We report them in [Appendix B.6, Supplementary Material](#). They confirm that people's wealth the first time we observe them is informative about

4. Our bootstrap procedure takes into account the sampling error arising from the limited number of observed individuals, but not from the limited number of periods for which each individual is observed. While the precision of the estimated fixed productivity increases with the latter, the length of observed periods per person is a common limitation of any panel survey dataset. To maximise the number of observations per person in our estimation, we utilise the long panel structure of the PSID and use twenty-four waves of data (from 1984 to 2017). In addition, when estimating our health process, we classify individuals into relatively large groups based on the terciles of the fixed productivity distribution. The latter also mitigates the sampling error concern due to the fact that we classify individuals in bins rather than using our point estimates for fixed productivity.

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

their 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 B3 in Appendix B.8, Supplementary Material](#). The top, middle, and bottom panels refer to the case in which people are currently in poor, fair, and good health, respectively. Different line styles represent 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 health 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 (solid lines in top left panel). If he is of the best health type η_3 he has a 20% probability of staying in poor health next period, while if he is of the worst health type η_1 , this probability increases to 80%. If, instead, he has been in bad health for three periods (lines with markers), 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.9, Supplementary Material](#) 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 about 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.

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 labour productivity, based on our estimates of equation (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 B7 in Appendix B.8, Supplementary Material](#).

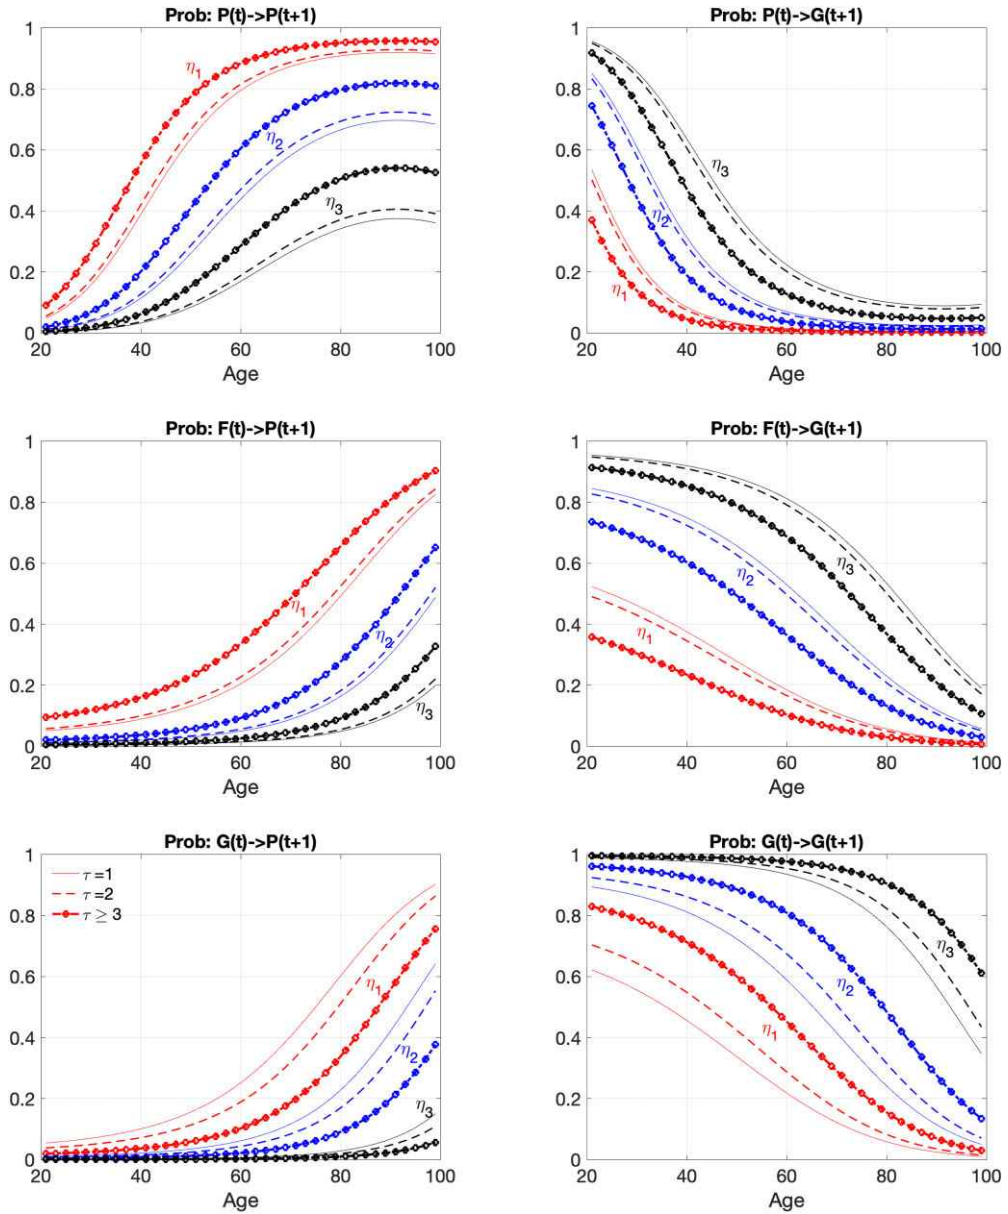


FIGURE 4

Estimated 2-year transition probabilities ($T = 3$)

Notes: The group of lines corresponding to each health type i is jointly labeled as η_i . Within each group, different line types correspond to the number of consecutive periods in current health status (τ): with the solid line referring to $\tau = 1$, the dashed line to $\tau = 2$, and the line with markers to $\tau \geq 3$

At age 21, only 8.3% of people are of the worst health type. The rows from second to fourth 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.

TABLE 3

Top row: Initial distribution of individuals over health types at age 21. Other rows, joint distribution of $\{\eta, \gamma\}$ at age 21

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

Notes: The terms γ_L , γ_M , and γ_H refer to the bottom, middle, and top terciles of the estimated fixed productivity distribution, respectively.




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

Notes: Bad health includes both fair and poor health

2.4.4. Comparing the implications of our processes with the data. We now turn to showing that our health and labour 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 2-year transition probabilities between good and bad health by age, as well as their duration dependence.

For an additional external validation of our health process, we turn to a balanced panel of males with a high-school degree aged 55–66 years old in the HRS. Figure 5 reports the number of waves spent in bad health over a 10-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 10 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 four to five waves during the next 10 years. The right panel shows that among those already unhealthy at age 55–56, 51% stay unhealthy during the next four to five 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.

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 labour productivity. Figure 6 displays the fraction of people in the bottom two terciles of fixed labour productivity (blue lines) or bottom tercile (red lines) among those in good health




FIGURE 6

Fraction of people in the bottom two terciles of fixed labour productivity ($\gamma_L + \gamma_M$) or bottom tercile of labour productivity (γ_L) among those in good health (left panel) and bad health (right panel)

Notes: 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

(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 labour productivity in our estimation, *i.e.* the health transitions are not a function of fixed labour productivity (see equations (2)–(5)). Thus, the ability of our model to reproduce it gives our model additional credibility.

2.5. What are the 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, several 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 10-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.

TABLE 4

Characteristics of a balanced panel of people who are healthy at age 55–56 and alive for the subsequent 10 years, organised by the number of unhealthy periods between ages 57 and 66

# Unhealthy periods	Individuals' characteristics (HRS)					Model	
	Smoking (%)	BMI	Father alive (%)	Mother alive (%)	Parents' educ (years)	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

Notes: 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.

TABLE 5

Average polygenic scores of people who are healthy at age 55–56 and alive for the subsequent 10 years, organised by the number of unhealthy periods between ages 57 and 66

# 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

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 10 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 four to five 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 (PGSs) are indices created by combining genetic markers that predict certain individuals' outcomes (for more details, see [Barth *et al.*, 2020](#); [Papageorge and Thom, 2020](#)).

We report four PGSs: for educational attainment, for lifestyle-related behaviours 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 score is normalised 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 four to five 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, Supplementary Material](#) shows that this finding is robust when we use the 25th, 50th, and 75th percentiles of the PGS 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 stylised 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, Supplementary Material](#).

3. WHAT ABOUT ENDOGENOUS HEALTH?

Our paper assumes that one's health evolution is pre-determined and thus exogenous in adulthood. That is, we do not allow either monetary investments or healthy behaviours to affect the evolution of health. A key reason for our choice is that we want to identify the effect of health types while keeping our model tractable, both in terms of computations and parameter identification. But because the evolution of health is an important topic, in this section, we discuss what we know from the literature in terms of the effects of medical spending on health, and what the PSID data tells us about the effects of exercise, a potentially key determinant of health (see for instance [Cole *et al.*, 2019](#)) for which the PSID has good data.

Starting from medical expenses, the key problem is that sick people spend more on medical goods and services and tend to stay sick longer and die faster than people who are healthy and spend less on medical goods and services. Thus, the raw data makes it hard to find that medical spending has a positive effect on health. Among the studies which use some credible (and hard to find) exogenous sources of variation in medical spending, many find that medical expenses have small or no effects on health and mortality (for instance [Brook *et al.*, 1983](#); [Fisher *et al.*, 2003](#); [Finkelstein and McKnight, 2008](#); [Black *et al.*, 2017](#)). Also, and importantly, [Danesh *et al.* \(2024\)](#) shows that in the Netherlands, where medical care is free, there is no differential access to medical care and yet people with higher income live substantially longer than those with lower income. This paper also finds that a large share of these mortality differences is explained by chronic diseases whose prevalence is already heterogeneous by income earlier in life. Hence, these findings indicate that there are important factors determining mortality and health differences that have little to do with the health care system and heterogeneity in medical spending. Given that it is difficult to convincingly establish to what extent medical spending improves health and extends lives, and that in many cases the empirical evidence suggests that these effects can be small, we abstract from this interesting force in this paper.

Turning to exercise, our goal is to evaluate whether allowing for a link between exercise and health overturns our findings about the importance of health types. Thus, we include exercise in our model of health dynamics.

More specifically, the PSID asks people about the number of times per week that they engage in light or heavy exercise. In each survey wave, we assign people who report doing no exercise to a "no exercise" category, those who report doing *only* light exercise at least once per week to "light" exercise, and those who report doing heavy exercise at least once per week (while also possibly engaging in light exercise) to "heavy" exercise.⁶ We denote with D_{EX} the dummy

6. There is a strong correlation between light and heavy exercise, hence people who engage in heavy exercise often also report engaging in light exercise.

TABLE 6
Estimation results for the health process including exercise behaviour

	$T = 5$	$T = 4$	$T = 3$	$T = 2$	$T = 1$
Coefficients of history-dependence terms and health types including exercise					
a_2^B	0.094	0.031	0.048	0.244	—
a_3^B	0.827**	0.656*	0.618**		
a_4^B	0.755**	0.630			
a_5^B	1.052***				
$a_{\eta_1}^B$	1.576***	1.679***	1.894***	2.149***	2.157***
$a_{\eta_3}^B$	-1.773***	-1.459***	-1.031	-1.202***	-1.395***
Light#poor health	0.324	0.250	0.158	0.030	0.059
Light#fair health	0.007	0.055	0.027	-0.002	0.026
Heavy#poor health	-0.248	-0.284	-0.328	-0.321	-0.339
Heavy#fair health	-0.505**	-0.530**	-0.554**	-0.606**	-0.569**
Coefficients of history-dependence terms and health types including exercise					
a_2^G	-0.380	-0.277	-0.383*	-0.535***	—
a_3^G	-0.0803	0.121	-0.741***		
a_4^G	-0.872**	-1.072***			
a_5^G	-1.597***				
$a_{\eta_1}^G$	2.509***	2.245***	1.961***	2.066***	2.023***
$a_{\eta_3}^G$	-1.596***	-1.991***	-2.516***	-2.647***	-2.815***
Light	-0.248	-0.224	-0.248	-0.298*	-0.328*
Heavy	-0.702***	-0.754***	-0.722***	-0.752***	-0.741***
N	8,799	9,521	10,381	11,279	12,181

Notes: The columns refer to specifications controlling for different number of lags of past health. The terms $a_{\eta_2}^B$ and $a_{\eta_2}^G$ are normalised to zero. No exercise, being in bad/good health for one period ($\tau_B = 1$, $\tau_G = 1$) and the first tercile of fixed productivity (γ_1) is the base case. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

for each exercise category. We modify our health transition dynamics as follows. We add the term $\mathbf{a}_E^G D_{EX}$ for the health transitions from good health in equations (4) and (5), and the term $\mathbf{a}_E^B D_{EX} D_{h_i}$ for the transitions from poor or fair health in equations (2) and (3). In the latter, D_{h_i} is the dummy variable of whether an individual is currently in poor or fair health. This way we allow for possibly different effect of exercise for those in fair versus poor health. Our health type prediction equation remains that in equation (6).

Furthermore, we estimate these processes by making the identifying assumption that a future health shock realised over the next 2 years (*i.e.* between t and $t + 1$) does not influence exercise in the current period. Cole *et al.* (2019) follow a similar strategy for 6-year health transitions.

Table 6 reports the resulting estimates for these health processes, while Appendix C.3, Supplementary Material displays those for our health types prediction equation. Several things are worth noticing about these results. First, both the size and the significance of the coefficients in our original specification without exercise are very consistent with the specification that includes exercise. In particular, the effect of health types remains important for one's future health evolution even when conditioning for exercise behaviour. Second, light exercise has little effect on one's future health dynamics, regardless of one's previous health history. Third, some heavy (and a combination of heavy and light) exercise does play some role in preventing future bad health, but only if one is in good or fair health to start with. Instead, once one is in poor health, both light and heavy exercise do little to improve one's future health outcomes. Fourth, and importantly, the coefficients on one's health types are larger than those on exercise, indicating that not only health types remain important once we condition for exercise, but they have a larger impact.

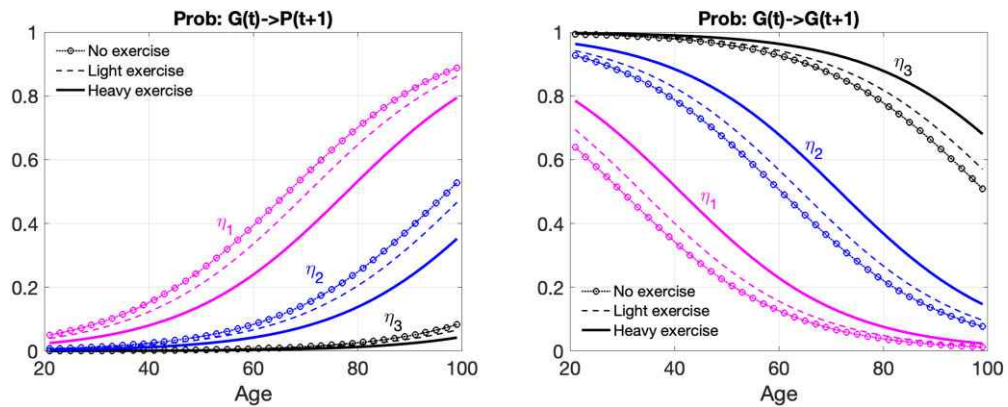


FIGURE 7

Estimated 2-year probabilities of turning to poor health (left panel) and good health (right panel) after having been in good health for one period (when keeping track of three periods of past health) by exercise category

Notes: On the left-hand side, the top three lines refer to the low health types (η_1), the middle three lines to the middle health type (η_2), and the bottom three lines to the best health types (η_3). On the right-hand side, the ordering of these groupings is reversed. Within each grouping, different line types correspond to exercise behaviour, with the dotted line referring to the no exercise group, the dashed line to the light exercise group, and the solid line to the heavy exercise group

To better illustrate the effects of each exercise category (none, light, heavy), Figure 7 reports a subset of transition probabilities for our baseline health process (in which we keep track of health during the previous three periods) for people who have been in good health for one period. The graph on the left-hand side plots the estimated probability of turning to poor health over the next 2 years, while the graph on the right-hand side plots the probability of remaining in good health.

The figure helps quantify what we already see in the tables: people who exercise have better health transitions compared to those who do not, and the difference is especially pronounced for heavy exercise. However, the variation in health transitions due to exercise is much smaller than the variation due to health types.

For instance, at age 61 the probability of turning to poor health for η_1 types (top three lines) decreases from 41% to 35%, and to 25% as one goes from no exercise to heavy exercise. The corresponding numbers for η_2 types (middle three lines) fall from 9% with no exercise to 4.5% with heavy exercise. For η_3 types (bottom three lines), exercise only marginally affects the chance of moving to poor health. Interestingly, the graph on the right-hand side shows that exercise increases the probability of staying in good health for η_3 types from 92% with no exercise to 96% with heavy exercise. When comparing across health types we can see that the probability of turning to poor health for a 61-year-old individual who does not exercise is 41% for η_1 , 7% for η_2 , and 0.4% for η_3 . The gaps by type in the probability of staying healthy are also big.

To quantify the maximum possible contribution of exercise to health over one's lifetime (rather than just from one period to the next) and compare it with that of health types, we next simulate a large number of individuals for each health type, and for whom we fix one's exercise behaviour to be always the same over all of the life-cycle. That is, every person can only exercise as follows: either never exercise, or always engage in light exercise, or always exercise heavily. We then use these simulated paths to compute the expected number of unhealthy periods between the ages of 21 and 80 by one's health type and exercise behaviour.

Table 7 reports the results and confirms that exercise does matter but that health types have a larger effect. For instance, comparing columns reveals that people of the worst health type (η_1) can decrease the average number of lifetime unhealthy periods from 14.4 to 12.3 (hence by over

TABLE 7

Average number of unhealthy periods (each lasts 2 years) between 21 and 80 years old by health type when keeping track of health during the current period and the two previous ones

Health type	Exercise category		
	<i>Always none</i>	<i>Always light</i>	<i>Always heavy</i>
η_1	14.4	13.9	12.3
η_2	7.0	6.4	4.5
η_3	1.0	0.8	0.5
All	5.7	5.3	4.0

Notes: The first, second, and third columns correspond to the cases when individuals never exercise, *Always* only do light exercise, or *Always* engage in heavy exercise.

4 years), by going from no exercise to always exercising heavily. People of the intermediate health type (η_2) can decrease the number of lifetime unhealthy periods from 7 to 4.5 (5 years), while those of the best health type (η_3) can reduce their lifetime unhealthy periods from 1 to 0.5 (1 year). However, comparing rows shows that among people who never exercise the number of lifetime unhealthy periods goes from 14.4 for the worst health types, to 1 for the best health types. The variation by types is similarly large for light and heavy exercise behaviour.

Thus, we find that even when controlling for exercise behaviour, health types are a key determinant of the number of unhealthy periods during one's life, which in turn, is a crucial determinant of the lifetimes costs of bad health. We now turn to our structural model to quantify the effects of these forces.

4. 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.

4.1. *Demographics, preferences, and labour income*

A model period is 2 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), Pashchenko and Porapakkarm (2016), Pashchenko and Porapakkarm (2017), and Rust and Phelan (1997).

Second, individuals differ in ex-ante characteristics, that is in their fixed labour productivity, health type, and patience, and these characteristics can be correlated with each other.⁷ This correlation captures the fact that some factors that are determined early in life can influence

7. Many studies find heterogeneity in patience, including Epper *et al.* (2020), Lawrance (1991), and Warner and Pleeter (2001). In addition, Cronqvist and Siegel (2015) find that genetic differences explain a significant fraction of

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 $m \in \{1, 2, 3\}$. The total number of types in the model is $3 \times 2 \times 3$, that is three health types, two patience types, and three fixed productivity types. We assume that β and γ are correlated through the health types. Hence, controlling for one's health type, β and γ are independent.

An individual is endowed with one unit of time that can be used for either leisure or work. Labour 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:

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

where ρ is risk aversion and \bar{n}_t is an age-specific household size.⁸ We follow Hall and Jones (2007) by adding a positive term \bar{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

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

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

$$z_t^h = \lambda_t^h \Upsilon_t, \quad (9)$$

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 4.4.

the variation in saving propensities across individuals, potentially through a link to the rate of time preferences or self-control. Several macroeconomic studies point out the importance of heterogeneity in rates of time preferences to explain wealth inequality (Krusell and Smith, 1998; Samwick, 1998; Hendricks, 2007a, 2007b; Krueger *et al.*, 2016; Carroll *et al.*, 2017).

8. We incorporate the family size \bar{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 5.1 for details on the construction of the targeted wealth profile). We take the average family size \bar{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.

4.2. Medical expenses and health insurance

During each period, every agent receives a medical expense shock (x_t^h) 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, Supplementary Material](#) 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 $cvg(x_t^h, i_H)$ 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 $cvg(x_t^h, 0) = 0$.

A working-age individual receives an offer to buy ESHI with probability $Prob_t$, which depends on age (t), fixed labour productivity (γ), and health ($h_t \in \{G, F, P\}$). We estimate $Prob_t$ from the MEPS. The variable g_t characterises 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 $p_I(h_t, t)$, 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:

$$p_I(h_t, t) = \zeta EM_t(h_t, t) + \varphi^h. \quad (10)$$

The term ζ 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

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

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

4.3. 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 \bar{y}_{ss} . There is a consumption tax τ_c and a tax on capital income τ_k . There is a progressive labour income tax $\mathcal{T}(y)$ which we specify as [Heathcote *et al.* \(2020\)](#)

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

The progressivity of the tax system is captured by $a_{\tau 1}$. We explain how we set taxes and transfers in [Appendix D.2, Supplementary Material](#).

We represent several existing means-tested programmes (Medicaid, food stamps, Disability Insurance, and Supplement Security Income) in a stylised way through a public safety-net programme, $T^{SI}(\bar{c})$. This programme guarantees every household a minimum consumption floor \bar{c} . This floor also captures the existence of uncompensated care or medical bankruptcy. In fact, in 2004, 85% 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 v_{R-1} (see Section 4.4), and on their ex-ante heterogeneity (γ, η, β) . Then, for each group, we compute average earnings over the seventeen 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 $ss(h_{R-1}, v_{R-1}, \gamma, \eta, \beta)$.

4.4. The labour productivity shock

Our labour productivity process is given by

$$z_t^h = \lambda_t^h \Upsilon_t = \lambda_t^h \exp(v_t) \exp(\gamma), \quad (12)$$

$$\begin{aligned} v_t &= \rho_v v_{t-1} + \varepsilon_t; & \varepsilon_t &\sim N(0, \sigma_\varepsilon^2), \\ \gamma &\sim N(0, \sigma_\gamma^2), \end{aligned} \quad (13)$$

where λ_t^h is a deterministic component that depends on age and health ($h_t \in \{G, F, P\}$) and the idiosyncratic component Υ_t consists of a persistent shock v_t and fixed productivity γ . We assume that γ is normally distributed and [Appendix B.5, Supplementary Material](#) shows that this assumption is consistent with the data. To obtain the initial distribution of fixed productivity and health type, we discretise γ into three terciles. Table 3 reports the conditional probability distribution from our maximum likelihood estimation.

To account for selection, we estimate the deterministic component of labour productivity λ_t^h in the second stage of our estimation procedure, as in [French \(2005\)](#). This is important 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. Though computationally costly, our strategy also ensures that the model reproduces the income-health gradient in the data (Figure 10 and Table 9), which is important for evaluating the costs of bad health through the labour market channel.⁹ [Appendix D.3, Supplementary Material](#) details the estimation of the stochastic part of the productivity component.

4.5. 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 labour supply (l_t) and insurance choice (i_H). At the end of the period, the medical expenses

9. An alternative method would be to perform the Heckman selection correction to the data, but this approach requires a variable serving as a valid exclusion restriction, which is typically hard to find.

shock (x_t^h) is realised. 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.

4.6. The optimisation 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 (v_t), ESHI offer status ($g_t \in \{0, 1\}$), age ($t \in \{1, 2, \dots, R - 1\}$), fixed productivity ($\gamma \in \{\gamma_L, \gamma_M, \gamma_H\}$), health type ($\eta \in \{\eta_1, \eta_2, \eta_3\}$), 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

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

where

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

subject to

$$\begin{aligned} k_t (1 + (1 - \tau_k) r) + z_t^h l_t - x_t^h (1 - cvg(x_t^h, i_H)) - P_t^h - Tax + T^{SI}(\bar{c}) \\ = (1 + \tau_c) c_t + k_{t+1} \end{aligned} \quad (16)$$

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

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

$$Tax = T(z_t^h l_t) + \tau_{MCR} z_t^h l_t + \tau_{ss} \min(z_t^h l_t, \bar{y}_{ss}) \quad (18)$$

$W_t(\mathbb{S}_t; l_t, i_H, x_t^h)$ is the interim value function conditional on the labour supply and insurance choices after the medical shock is realised. The conditional expectation on the right-hand side of equation (15) is over $\{h_{t+1}, z_{t+1}^h, g_{t+1}\}$. 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 (x_t^h), health status at 64 ($h_{R-1} \in \{P, F, G\}$), productivity shock before retirement (v_{R-1}), age ($t \in \{R, \dots, T\}$), fixed productivity 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

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

where

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

subject to

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

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

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

5. 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.

5.1. 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, Supplementary Material](#). The survival probability and health 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 minimise the weighted sum of square differences 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 $\{\bar{b}, \phi_W, \phi_F, \phi_P, \lambda_t^G, \lambda_t^F, \lambda_t^P, \beta_{low}, \beta_{high}, Pr(\beta_{low}|\eta), Pr(\beta_{high}|\eta), \theta_{Beq}, k_{Beq}, \bar{c}\}$. 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 corresponding 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:¹⁰

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

[Viscusi \(1993\)](#) provided an extensive review documenting that the estimates vary from \$1 million to \$16 millions (in 1990 dollars). The U.S. government agencies (Department of Transportation,

10. 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.

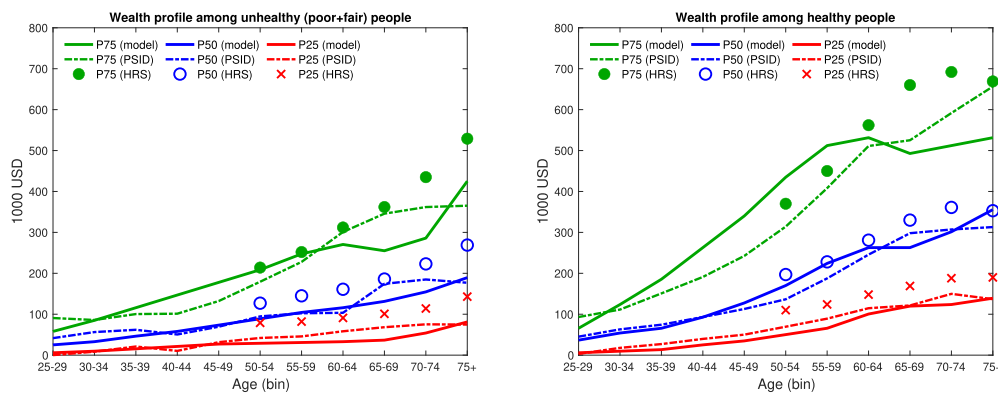


FIGURE 8

Wealth profiles by health status: data versus model

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 parameterisation the non-pecuniary effects of health are large. In Appendix E.4, Supplementary Material, we report the results from an alternative parameterisation when the VSL is equal to \$6 millions, which emphasises the non-pecuniary effects of health more.

Labour market outcomes. We target the fraction of workers and average labour income conditional on working for each age and health status (good, fair, and poor). Figure 10 reports our targeted moments from the PSID as dots/crosses and our model implications as solid lines. The labour income profiles in the right panel come from our estimated coefficients in equation (1) and include fixed productivity (γ) for each age and health. Since the estimated labour income profiles among workers in poor and fair health are quite similar, we estimate their targeted labour income profiles from the pooled sample of people in poor and fair health. Related, it is worth pointing out that, even though the targeted labour income profile for people in poor and fair health is the same, their estimated deterministic productivities (λ_i^F , λ_i^P) are different due to their selection into employment.

Wealth moments. We target the 25th, 50th, and 75th percentiles of wealth, conditional on being healthy and unhealthy (poor + fair) by 5-year age windows (dashed lines in Figure 8). 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).¹¹ Because net worth is measured at the household level and our model abstracts from heterogeneity in family size, we adjust observed wealth by family size as follows:

$$wealth_{it} = \sum_{j=G,B} \left(d_{age}^j D_{it}^{age} + d_1^j n_{it} + d_2^j n_{it}^2 + d_3^j n_{it}^3 \right) \mathbf{1}_{\{h_{it}=j\}} + \sum_{t=1994}^{2013} d_t D_t + res_{it}, \quad (22)$$

11. 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 convert it to 2013 dollars using the CPI.




FIGURE 9

Wealth gradient: data versus model

where $wealth_{it}$ is net worth, D_{it}^{age} 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, \bar{n}_t , to get our measure of net worth.¹² Then we construct the targeted 25th, 50th, and 75th percentiles of wealth distribution among people in good and bad health, and report them as dashed lines in Figure 8. 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 8.

The wealth profiles from the two datasets are remarkably similar. Figure 8 displays the wealth–health gradient typically documented in the literature: Figure 9 emphasises 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.

5.2. Second step estimation results

The third column of Table 8 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.¹³ 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 decumulation 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

12. We compute $\sum_{j=G,B}(\hat{a}_{age}^j D_{it}^{age} + \hat{a}_1^j \bar{n}_t + \hat{a}_2^j \bar{n}_{age}^2 + \hat{a}_3^j \bar{n}_{age}^3) \mathbf{1}_{\{h_{it}=j\}} + \hat{a}_{2013} + \widehat{res}_{it}$, where \hat{a}^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.

13. In Appendix E.1, Supplementary Material, 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.

TABLE 8
Preference parameters and the consumption floor

Parameters		Baseline	No correlation
(Annual) discount factor	$\{\beta_{low}, \beta_{high}\}$	{0.877, 0.992}	{0.895, 0.992}
% β_{low} by η_i at age 20	$Pr(\beta_{low} \eta_1)$	78%	55.6%
	$Pr(\beta_{low} \eta_2)$	79%	55.6%
	$Pr(\beta_{low} \eta_3)$	38%	55.6%
Bequest parameter	θ_{Beq}	1,905	1,256
	κ_{Beq}	\$182,707	\$168,577
(Annual) consumption floor	\bar{c}	\$3,505	\$4,116

Notes: Both β and \bar{c} are converted into annual values.

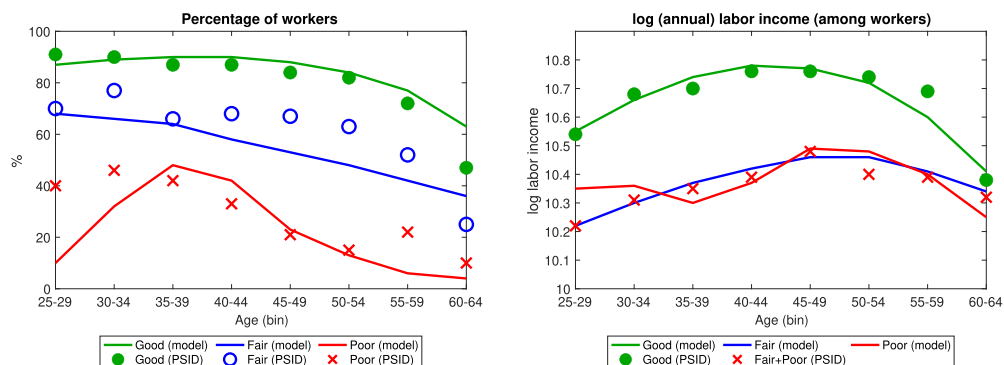


FIGURE 10

Employment by health (left panel) and average labour income among workers by health (right panel)

Notes: The dotted and crossed marks are from the PSID while the solid lines are from our model

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 labour supply. More specifically, Capatina's (2015) estimate of the consumption floor is \$4,114 (in 2006 USD) while Pashchenko and Porapakarm's (2017) estimate is \$1,540 (in 2003 USD).

5.3. Model fit

Figure 10 compares the employment rate (left panel) and the average labour income of workers (right panel) generated by our model (solid lines) with the targeted profiles from the PSID (dots and crosses).¹⁴ Our model matches the important differences in labour market outcomes across health status very well.

14. The annual labour income profiles from the model are constructed from dividing 2-year labour income by two.

TABLE 9
Percentage of unhealthy (poor + fair) individuals in each earnings tercile

	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

Notes: Left panel: data from PSID (data from HRS is in parentheses). Right panel: model.

TABLE 10
Percentage of unhealthy (poor + fair) individuals in each wealth tercile

	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

Notes: Left panel: data from PSID (data from HRS is in parentheses). Right panel: model.

Figure 8 displays the wealth profiles from our model (solid line) and the data (with dashed lines for the PSID and dots/crosses for the HRS).¹⁵ Our model matches the wealth gap between healthy and unhealthy people for the 25th, 50th, and 75th 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 9 and 10 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.

6. 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

15. 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; 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 modelling the saving decisions of the very rich, we abstract from these mechanisms.

TABLE 11
Unconditional wealth quartiles at age 60–64 (in 1,000 USD).

Wealth percentile	PSID (HRS)	Baseline	No correlation $Pr(\beta_{low} \eta_i) = 0.556$
25th pct	\$92 (\$117)	\$89	\$93
50th pct	\$213 (\$220)	\$224	\$201
75th pct	\$476 (\$383)	\$502	\$435

measure of the monetary costs of bad health over the life-cycle. Third, we evaluate the welfare losses due to bad health realisations, 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.

6.1. *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 quantify its importance, we estimate an alternative model in which one’s discount factor is orthogonal to one’s health type and, consequently, to one’s health status.

Formally, we set $Pr(\beta_{low}|\eta_m) = 0.556$ for all health types. This number corresponds to the overall fraction of people with low patience in our baseline economy. Then, we re-estimate our model by matching the same targets. We call this version of our model the “no-correlation” model. It is important to notice that it still features all of the channels through which bad health can affect individuals’ savings.

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 8 for their values). 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, 2007a, 2007b). Our model-simulated data matches the wealth quartiles unconditional on health status. See the second and fourth columns of Table 11 for a comparison of the 25th, 50th, and 75th 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 labour income by health status conditional on working (the income–health gradient).

The “no-correlation” model, however, falls short of replicating the observed large differences in wealth by health status (see Table 12). For example, for people near retirement, the difference between the median wealth of the healthy and that of the unhealthy is only \$35k in this no-correlation model, compared to about \$140k in the PSID and our 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

TABLE 12
Wealth–health gradient at age 60–64 (in 1,000 USD)

Wealth difference by health status	PSID (HRS)	Baseline	No correlation $Pr(\beta_{low} \eta_i) = 0.556$
25th pct	\$56 (\$47)	\$67	\$58
50th pct	\$142 (\$98)	\$146	\$35
75th pct	\$210 (\$222)	\$260	\$99

Notes: The table reports the wealth difference between healthy and unhealthy (poor + fair) people for each wealth quartile.

(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.

6.2. The monetary losses due to bad health

While previous literature has noted that health deterioration leads to worse outcomes, a comprehensive evaluation of the effects of bad health was not previously done. To achieve this goal, we fold the consequences of bad health on a number of dimensions (income, medical spending, etc.) into one measure. We do so by comparing each individual with his counterfactual self when he does not face any bad health realisation throughout his lifetime, but faces the same environment. This computation allows us to capture both the pecuniary (by comparing net income) and non-pecuniary (by comparing utility) costs of bad health over all of the life-cycle. In addition, it allows us to decompose the portion of health costs that is due to ex-ante differences (the types) and ex-post bad luck (the shocks).

Importantly, our approach has an important advantage over estimating the costs of bad health directly from the data. Without a model, in fact, we can only compare *different* individuals, either healthy or unhealthy, after controlling for observable characteristics. But, as we show earlier in this paper, healthy and unhealthy people differ in unobservable characteristics (health type, permanent productivity, and preferences), and this biases the estimated costs of bad health. Instead, we compare the exact same people, in the exact same environment.

To formalise our computation, 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 $\frac{1}{\hat{T}} \sum_{t=1}^{\hat{T}} \frac{y_{it}^H - y_{it}^{BS}}{(1+r)^t}$, where \hat{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, [Supplementary Material](#).

Table 13 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 are worth noticing. First, on average, because of bad health realisations, 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.

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

TABLE 13
Annual monetary losses due to bad health (poor + fair)

	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)	(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)	(2.7%)	(18%)	(3%)	(0.2%)

Notes: 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%. Average earning in our baseline model is \$38,648 per year.

TABLE 14
Composition of annual monetary loss due to bad health (poor + fair) using 2% interest rate when computing the present value

	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

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.

Next, we turn to decomposing the sources of the monetary losses due to bad health. Table 14 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 21–64). 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.

TABLE 15
Average age at death in the baseline case

	Baseline					When everyone is always healthy
	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

Notes: In the counterfactual scenario, since everyone is healthy, the distribution of age at death is the same for all η_i .

6.3. 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 15 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.

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 realised lifetime utility of an individual in the baseline is

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

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 equals 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 (realised) 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:

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

where $\{c_t^{**}, l_t^{**}, k_t^{**}\}$ 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 $\lambda_c \bar{c}^{**}$, where \bar{c}^{**} is one's average consumption over the lifetime when drawing only good health realisations.

Table 16 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

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.

TABLE 16
Welfare losses due to bad health (poor + fair)

	all	η_1	η_2	η_3	β_L	β_H
Compensated consumption equivalence (% consumption equivalence, λ_c)	\$1,933 (10.6)	\$6,380 (36.8)	\$2,690 (14.8)	\$854 (4.4)	\$1,866 (10.3)	\$2,018 (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

Notes: The dollar value is calculated from $\lambda_c \bar{c}^{**}$ where λ_c and \bar{c}^{**} are the percentage of consumption reduction and average life time consumption of each individual when always healthy.

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 realisations 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 health type spend more than half of their life in bad health. Table E13 in Appendix E.3, Supplementary Material reports welfare losses by patience, health, and productivity types.

The other lines of Table 16 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 CCE by comparing one of the counterfactual baselines with the hypothetical situation of no bad health realisations. 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 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 17 display

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.

TABLE 17
Concentration of losses due to bad health (poor + fair) and variation of losses due to health types

	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

Notes: The reported numbers in columns 2–4 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 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).

The last column of Table 17 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 13), 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 15). 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, is due to fixed health types and is thus related to factors that are likely pre-determined earlier in life.

7. CONCLUSIONS AND DIRECTIONS FOR FUTURE RESEARCH

Understanding and quantifying an issue is a necessary condition to design an effective policy intervention. We develop a structural framework for evaluating the differences in economic outcomes by health and for measuring the lifetime consequences of being unhealthy over the life-cycle. Our approach emphasises the complex nature of health dynamics and the role played by multidimensional ex-ante differences across individuals. Our measurement exercise provides a comprehensive assessment that can be used as a starting point for policy analysis.

We provide several important findings. Regarding our estimated health 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 ex-ante 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 labour 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 (in present terms) over the life-course because of bad health realisations. In welfare terms, bad health realisations 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 labour 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 health is large, especially for monetary costs.

Thus, our measure of the lifetime costs of bad health and its decomposition emphasises several points that are important to take into account when designing policy. First, that bad health creates large costs that accumulate over one's lifetime. Second, that these costs far exceed those of medical expenses (in fact, we find that foregone labour income and lower life expectancy are key determinants of the costs of bad health). Third, that, to a significant extent, one's lifetime costs of bad health are pre-determined by one's health type, *i.e.* by genetic endowments and early life circumstances.

Thus, our measure of the lifetime costs of bad health and its decomposition emphasises several points that are important to take into account when designing policy. First, bad health creates large costs that accumulate over one's lifetime. Second, these costs far exceed those of medical expenses (in fact, we find that foregone labour income and lower life expectancy are key determinants of the costs of bad health). Third, to a significant extent, one's lifetime costs of bad health are pre-determined by one's health type, *i.e.* by genetic endowments and early life circumstances.

These findings have several important policy implications. To start, health insurance only partly insures the costs of bad health in that it covers, at best, only medical expenses, which are only one portion of the costs of bad health. Then, labour market policies targeted towards unhealthy people are potentially very beneficial. Finally, most effective policies should target one's health type's formation. These policies include various genetic treatments and improving a child's early life environment. Our results thus move the focus to these policies and our paper provides a framework that can be used to better measure the long-term consequences of improving people's health types earlier during the life-cycle.

A second important direction for future research is to better understand the direct effects of health on labour market productivity. Our results show that an important part of the monetary costs of bad health comes from its effect on labour income. It is possible that this occurs because health affects human capital accumulation. Modelling 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 labour productivity? In our approach, we estimate fixed labour 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 ex-ante difference across individuals. Our estimates show that fixed characteristics, namely, health types, patience, and fixed labour 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 modelling these issues requires incorporating the childhood stage of the life-cycle and modelling 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 emphasise the importance of fixed characteristics for generating large losses due to bad health. Hence, developing policies to prevent these losses requires understanding the formation of these characteristics.

Acknowledgments. We thank Liran Einav, Christopher Flinn, Eric French, Joseph Altonji, Selahattin Imrohorglu, John Kennan, Naoki Aizawa, Sagiri Kitao, Rasmus Lentz, Vincenzo Quadrini, Yongseok Shin, Matthew White, and seminar participants at many institutions for comments and suggestions. D.N. gratefully acknowledges support from the ERC, grant 614328 “Savings and Risks.” P.P. gratefully acknowledges support from the JSPS KAKENHI Grant Number 15K03505. S.P. 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.

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online.

Data Availability Statement

The data and code underlying this research are available on Zenodo at <https://doi.org/10.5281/zenodo.10891439>.

REFERENCES

- ALMLUND, M., DUCKWORTH, A. L. and HECKMAN, J., *et al.* (2011), “Personality Psychology and Economics”, in Hanushek, E. A., Machin, S., and Woessmann, L. (eds) *Handbook of the Economics of Education* (Vol. 4) (Amsterdam: Elsevier) 1–181.
- AMERIKS, J., BRIGGS, J., CAPLIN, A., *et al.* (2020), “Long-Term Care Utility and Late in Life Saving”, *Journal of Political Economy*, **128**, 2375–2451.
- ANDA, R., FELITTI, V., BREMNER, D., *et al.* (2006), “The Enduring Effects of Abuse and Related Adverse Experiences in Childhood. A Convergence of Evidence from Neurobiology and Epidemiology”, *European Archives of Psychiatry and Clinical Neuroscience*, **256**, 174–186.
- ARTHUR, B. (1981), “The Economics of Risks to Life”, *American Economic Review*, **71**, 54–61.
- ATTANASIO, O., KITAO, S. and VIOLANTE, G. (2011), “Financing Medicare: A General Equilibrium Analysis”, in Shoven, J. B. (ed) *Demography and the Economy* (Cambridge, MA: National Bureau of Economic Research, Inc.) 333–366.
- BARTH, D., PAPAGEORGE, N. and THOM, K. (2020), “Genetic Endowments and Wealth Inequality”, *Journal of Political Economy*, **128**, 1474–1522.
- BLACK, B., ESPIN-SACHEZ, J., FRENCH, E., *et al.* (2017), “The Long-Term Effect of Health Insurance on Near-Elderly Health and Mortality”, *American Journal of Health Economics*, **3**, 281–311.
- BLUNDELL, R., BORELLA, M. and COMMAULT, J. (2020), “Why Does Consumption Fluctuates at Old Age and How Should the Government Insure It?” (Working Paper 27348, NBER).
- BLUNDELL, R., BRITTON, J. and COSTA DIAS, M., *et al.* (2023), “The Impact of Health on Labor Supply Near Retirement”, *Journal of Human Resources*, **58**, 282–334.
- BOLT, U. (2021), “What is the Source of the Health Gradient? The Case of Obesity” (Mimeo, UCL).
- BROOK, R., WARE, J. JR., ROGERS, W., *et al.* (1983), “Does Free Care Improve Adults’ Health? Results from a Randomized Controlled Trial”, *New England Journal of Medicine*, **309**, 1426–34.
- CAGETTI, M. and DE NARDI, M. (2006), “Entrepreneurship, Frictions, and Wealth”, *Journal of Political Economy*, **114**, 835–870.
- CAPATINA, E. (2015), “Life-Cycle Effects of Health Risk”, *Journal of Monetary Economics*, **74**, 67–88.

- CAPATINA, E., KEANE, M. and MARUYAMA, S. (2021), "Health and Earning Dynamics over the Life-Cycle" (Mimeo, University of New South Wales).
- CARROLL, C., SLACALEK, J., TOKUOKA, K., *et al.* (2017), "The Distribution of Wealth and the Marginal Propensity to Consume", *Quantitative Economics*, **8**, 977–1020.
- CASE, A., FERTIG, A. and PAXSON, C. (2005), "The Lasting Impact of Childhood Health and Circumstance", *Journal of Health Economics*, **24**, 365–389.
- CASE, A., LUBOTSKY, D. and PAXSON, C. (2002), "Economic Status and Health in Childhood: The Origins of the Gradient", *American Economic Review*, **92**, 1308–1664.
- CASE, A. and PAXSON, C. (2010), "Causes and Consequences of Early-Life Health", *Demography*, **47**, S65–S85.
- COLE, H., KIM, S. and KRUEGER, D. (2019), "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*, **56**, 1123–1169.
- CONESA, J., COSTA, D., KAMALI, P., *et al.* (2018), "Macroeconomic Effects of Medicare", *Journal of the Economics of Ageing*, **11**, 27–40.
- CONTI, G. and 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**, 585–605.
- CRONQVIST, H. and SIEGEL, S. (2015), "The Origins of Savings Behavior", *Journal of Political Economy*, **123**, 123–169.
- CURRIE, J. and MADRIAN, B. (1999), "Health, Health Insurance, and the Labor Market", in Ashenfelter, O. and Card, D. (eds) *Handbook of Labor Economics* (Vol. 3) (Amsterdam: Elsevier) 3309–3416.
- 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* (New York: Oxford University Press) 124–163.
- DANESH, K., KOLSTAD, J. and PARKER, W., *et al.* (2024), "The Chronic Disease Index: Analyzing Health Inequalities Over the Lifecycle" (Working Paper 32577, NBER).
- DE NARDI, M. (2004), "Wealth Inequality and Intergenerational Links", *Review of Economic Studies*, **71**, 743–768.
- DE NARDI, M., FRENCH, E. and JONES, J. (2010), "Why Do the Elderly Save?", *Journal of Political Economy*, **118**, 39–75.
- DE NARDI, M., FRENCH, E. and JONES, J. (2016), "Medicaid Insurance in Old Age", *American Economic Review*, **106**, 3480–3520.
- DE NARDI, M., PASHCHENKO, S. and PORAPAKKARM, P. (2017), "The Lifetime Costs of Bad Health" (Working paper 2017-079, HCEO).
- EPPER, T., FEHR, E. and FEHR-DUDA, H., *et al.* (2020), "Time Discounting and Wealth Inequality", *American Economic Review*, **110**, 1177–1205.
- FANG, H. and KRUEGER, D. (2022), "The Affordable Care Act after a Decade: Its Impact on the Labor Market and the Macroeconomy", *Annual Reviews of Economics*, **4**, 453–494.
- FELITTI, V., WILLIAMSON, D. and SPITZ, A., *et al.* (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**, 245–258.
- FINKELSTEIN, A. and MCKINIGHT, R. (2008), "What did Medicare do? The Initial Impact of Medicare on Mortality and out of Pocket Medical Spending", *Journal of Public Economics*, **92**, 1644–1668.
- FISHER, E., WENBERG, D., STUKEL, T., *et al.* (2003), "The Implications of Regional Variations in Medicare Spending. Part 1: The Content, Quality, and Accessibility of Care", *Annals Internal Medicine*, **138**, 273–287.
- FRENCH, E. (2005), "The Effects of Health, Wealth, and Wages on Labor Supply and Retirement Behavior", *Review of Economic Studies*, **72**, 395–427.
- FRENCH, E. and JONES, J. (2011), "The Effects of Health Insurance and Self-Insurance on Retirement Behavior", *Econometrica*, **79**, 693–732.
- GUVENEN, F. (2009), "An Empirical Investigation of Labor Income Processes", *Review of Economic Dynamic*, **12**, 58–79.
- HAI, R. and HECKMAN, J. (2021), "A Dynamic Model of Health, Addiction, Education, and Wealth" (Mimeo, University of Miami).
- HALL, R. and JONES, C. (2007), "The Value of Life and the Rise in Health Spending", *Quarterly Journal of Economics*, **122**, 39–72.
- HALLIDAY, T. (2008), "Heterogeneity, State Dependence and Health", *Econometrics Journal*, **11**, 499–516.
- HARRIS, S. E., HAGENARS, S. P. and DAVIES, G., *et al.* (2017), "Molecular Genetic Contributions to Self-Rated Health", *International Journal of Epidemiology*, **46**, 994–1069.
- HEATHCOTE, J., STORESLETTEN, K. and VIOLANTE, G. (2020), "Optimal Progressivity with Age-Dependent Taxation", *Journal of Public Economics*, **189**, 1040–1074.
- HENDRICKS, L. (2007a), "Retirement Wealth and Lifetime Earnings", *International Economic Review*, **48**, 421–456.
- HENDRICKS, L. (2007b), "How Important is Discount Rate Heterogeneity for Wealth Inequality?", *Journal of Economic Dynamics & Control*, **31**, 3042–3068.
- HOSSEINI, R., KOPECKY, K. and ZHAO, K. (2021), "How Important is Health Inequality for Lifetime Earnings Inequality?" (Mimeo, Federal Reserve Bank of Atlanta).
- HOSSEINI, R., KOPECKY, K. and ZHAO, K. (2022), "The Evolution of Health over the Life Cycle", *Review of Economic Dynamics*, **45**, 237–263.

- HUBBARD, G., SKINNER, J. and ZELDES, S. (1994), "The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving", *Carnegie-Rochester Conference Series on Public Policy*, **40**, 59–125.
- IDLER, E. and BENYAMINI, Y. (1997), "Self-Rated Health and Mortality: A Review of Twenty-Seven Community Studies", *Journal of Health and Social Behavior*, **38**, 21–37.
- JESKE, K. and KITAO, S. (2009), "U.S. Tax Policy and Health Insurance Demand: Can a Regressive Policy Improve Welfare?", *Journal of Monetary Economics*, **56**, 210–221.
- JUNG, J. and TRAN, C. (2016), "Market Inefficiency, Insurance Mandate and Welfare: U.S Health Care Reform 2010", *Review of Economic Dynamics*, **20**, 132–159.
- 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>.
- KRUEGER, D., MITMAN, K. and PERRI, F. (2016), "Macroeconomics and Household Heterogeneity", in Taylor, J. and Uhlig, H. (eds) *Handbook of Macroeconomics* (Vol. 2) (Chicago: Elsevier) 843–921.
- KRUSSELL, P. and SMITH, A. (1998), "Income and Wealth Heterogeneity in the Macroeconomy", *Journal of Political Economy*, **106**, 867–896.
- LANGE, F. and MCKEE, D. (2012), "The Evolution of Latent Health over the Life Course" (Mimeo, Yale University).
- LAWRANCE, E. (1991), "Poverty and the Rate of Time Preference: Evidence from Panel Data", *Journal of Political Economy*, **99**, 54–77.
- LOCKWOOD, L. (2018), "Incidental Bequests and the Choice to Self-Insure Late Life Risks", *American Economic Review*, **108**, 2513–2550.
- LOW, H. and PISTAFERRI, L. (2015), "Disability Insurance and the Dynamics of the Incentive Insurance Trade-Off", *American Economic Review*, **105**, 2986–3029.
- MURPHY, K. and TOPEL, R. (2006), "The Value of Health and Longevity", *Journal of Political Economy*, **114**, 871–904.
- NAKAJIMA, M. and TELYUKOVA, I. (2020), "Home Equity in Retirement", *International Economic Review*, **61**, 573–616.
- OZKAN, S. (2023), "Income Differences and Health Disparities: Role of Preventive vs. Curative Healthcare" (Working Paper 2023-025C, Federal Reserve Bank of St. Louis).
- PAPAGEORGE, N. and THOM, K. (2020), "Genes, Education, and Labor Market Outcomes: Evidence from the Health and Retirement Study", *Journal of the European Economic Association*, **18**, 1351–1399.
- PASHCHENKO, S. (2013), "Accounting for Non-Annuity", *Journal of Public Economics*, **98**, 53–67.
- PASHCHENKO, S. and PORAPAKKARM, P. (2013), "Quantitative Analysis of Health Insurance Reform: Separating Regulation from Redistribution", *Review of Economic Dynamics*, **16**, 383–404.
- PASHCHENKO, S. and PORAPAKKARM, P. (2016), "Cross-subsidization in Employer-Based Health Insurance and the Effects of Tax Subsidy Reform", *National Tax Journal*, **69**, 583–612.
- PASHCHENKO, S. and PORAPAKKARM, P. (2017), "Work Incentives of Medicaid Beneficiaries and the Role of Asset Testing", *International Economic Review*, **58**, 1117–1154.
- PASHCHENKO, S. and PORAPAKKARM, P. (2019), "Reducing Medical Spending of the Publicly Insured: The Case for a Cash-out Option", *American Economic Journal: Economic Policy*, **11**, 390–426.
- PIJOAN-MAS, J. and RÍOS RULL, V. (2014), "Heterogeneity in Expected Longevities", *Demography*, **51**, 2075–2102.
- POTERBA, J., VENTI, S. and WISE, D. (2017), "The Asset Cost of Poor Health", *The Journal of the Economics of Ageing*, **9**, 172–184.
- PRADOS, M. (2018), "Health and Earnings Inequality over the Life Cycle and The Redistributive Potential of Health Policies" (Mimeo, University of Southern California).
- QUADRINI, V. (1999), "The Importance of Entrepreneurship for Wealth Concentration and Mobility", *Review of Income and Wealth*, **45**, 1–19.
- ROBINSON, L. (2007), "How US Government Agencies Value Mortality Risk Reductions", *Review of Environmental Economics and Policy*, **1**, 283–299.
- ROSEN, S. (1988), "The Value of Changes in Life Expectancy", *Journal of Risk and Uncertainty*, **1**, 285–304.
- RUST, J. and PHELAN, C. (1997), "How Social Security and Medicare Affect Retirement Behavior In a World of Incomplete Markets", *Econometrica*, **65**, 781–831.
- SAMWICK, A. (1998), "Discount Rate Heterogeneity and Social Security Reform", *Journal of Development Economics*, **57**, 117–146.
- SHEPARD, D. and ZECKHAUSER, R. (1984), "Survival versus Consumption", *Management Science*, **30**, 423–439.
- 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**, 1621–1629.
- VISCUSI, K. (1993), "The Value of Risks to Life and Health", *Journal of Economic Literature*, **31**, 1912–1946.
- WARNER, J. and PLEETER, S. (2001), "The Personal Discount Rate: Evidence from Military Downsizing Programs", *American Economic Review*, **91**, 33–53.

Online Appendix: The Lifetime Costs of Bad Health*

Mariacristina De Nardi Svetlana Pashchenko Ponpoje Porapakkarm

November 9, 2023

*De Nardi: University of Minnesota, Federal Reserve Bank of Minneapolis, CEPR, and NBER; denardim@nber.org. Pashchenko: University of Georgia; svetlana@uga.edu. Porapakkarm: National Graduate Institute for Policy Studies; p-porapakkarm@grips.ac.jp.

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 seminar participants at many institutions for 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.

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.¹ 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 process, we further drop every other wave prior to 1997.

In addition to using the PSID to estimate the health 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 5.1, 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 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 HRS Longitudinal file 2016 (v1).

We use MEPS for our first step estimation. Specifically, we use it to estimate total medical

¹ Our results are robust to using equal weights.

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.² We use waves 1999/2000-2016/2017 of MEPS.

B 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 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 process in terms of the distribution of people by length of unhealthy spells.

B.1 Constructing 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:

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

² Pashchenko and Porapakarm (2016b) provide more details on the MEPS dataset.

B.2 Likelihood function for the two-year health 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 $t_0 + J$: $\{h_{i,t_0}, h_{i,t_0+1}, \dots, h_{i,t_0+J}\}$. If an individual is of η_j -type, the conditional probability of the observed health sequence can be constructed from Eq. (2)-(5) and the estimated conditional survival probability in Eq. (7). Denote the conditional probability of this health sequence as $Pr(h_{t_0+\tau}, \dots, h_{t_0+J} \mid h_{i,t_0}, \dots, h_{i,t_0+\tau-1}, \eta_j)$.

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

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

where $Pr(\eta_j \mid \mathbf{X}_{i,t_0}^\eta)$ 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,

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

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

B.3 Estimating 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 Using bootstrapping to estimate confidence intervals

In the main text, we proceed in two steps when estimating our health process in Eq.(2)-(6). We first estimate labor productivity from the PSID and survival probability from the HRS, and then use these estimates in our maximum likelihood estimation. In this section, we account for the estimation error in fixed labor productivity and survival probabilities by using bootstrapping to compute the 90 and 95% confidence intervals for our estimated health process and health types parameters. More specifically, we re-sample from both our HRS and PSID samples for 1000 times. For each of these draws, we re-estimate fixed productivity and survival probabilities, and use each set of results when estimating our health process with a maximum likelihood.

Tables B1 and B2 report our estimated coefficients and their 95% confidence interval (in parentheses and below each estimate). We mark each estimate with ** and * when the 95% and 90% bootstrapped confidence interval for that parameter excludes zero, respectively. This check shows that the inference drawn in this case is consistent with the one in our baseline case.

	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 (-0.31, 0.45)	0.0715 (-0.31, 0.42)	0.129 (-0.25, 0.52)	0.288 (-0.1, 0.74)	—
a_3^B	0.826** (0.21, 1.36)	0.810** (0.2, 1.32)	0.675** (0.14, 1.22)		
a_4^B	0.528 (-0.12, 1.08)	0.704* (-0.05, 1.21)			
a_5^B	0.772* (-0.03, 1.38)				
$a_{\eta_1}^B$	2.270** (1.35, 3.29)	1.994** (1.01, 3.1)	1.604** (1.09, 2.7)	1.987** (1.16, 2.58)	2.111** (1.69, 2.52)
$a_{\eta_3}^B$	-2.043** (-2.78, -0.77)	-2.104** (-3.08, -0.92)	-1.346** (-2.89, -0.0)	-1.303* (-2.21, -0.42)	-1.506** (-2.41, -0.44)
Coefficients of history-dependence terms and health types in Eq. (4)-(5)					
a_2^G	-0.391** (-0.68, -0.05)	-0.366** (-0.65, -0.001)	-0.369** (-0.7, -0.07)	-0.770** (-1.13, -0.48)	—
a_3^G	-0.242 (-0.58, 0.23)	-0.183 (-0.54, 0.33)	-1.086** (-1.7, -0.76)		
a_4^G	-1.007** (-1.55, -0.41)	-1.691** (-2.1, -0.92)			
a_5^G	-1.921** (-2.24, -1.36)				
$a_{\eta_1}^G$	4.527** (1.78, 26.2)	3.784** (1.79, 5.84)	1.637** (1.2, 5.7)	1.806** (1.24, 5.0)	2.006** (1.52, 2.54)
$a_{\eta_3}^G$	-1.447** (-1.98, -0.97)	-1.639** (-2.46, -1.1)	-2.318** (-3.07, -1.8)	-2.555** (-3.22, -2.24)	-2.871** (-3.39, -2.58)
N	9028	9765	11126	12096	13083

Table B1: Estimation results for the health process in Eq. (2)-(3) in the top panel and Eq. (4)-(5) in the bottom panel. The columns refer to specifications controlling for different number of lags of past health. The terms $a_{\eta_2}^B$ and $a_{\eta_2}^G$ are normalized to zero. Being in bad/good health for one period ($\tau_B = 1, \tau_G = 1$) is the base case. All estimations include a quadratic in age whose coefficients depend on current health status (poor, fair, good). The numbers in parentheses are the 95th confidence interval from bootstrapped samples. The asterisks ** (*) refer to estimates for which the 95% (90%) confidence interval excludes zero.

	T=5	T=4	T=3	T=2	T=1
age t_0	-0.039 (-0.15, 0.02)	-0.025 (-0.11, 0.02)	0.005 (-0.03, 0.05)	0.001 (-0.04, 0.03)	-0.016 (-0.05, 0.01)
$h_{t_0} = G$	-1.457** (-2.4, -0.81)	-1.429** (-2.32, -0.58)	-1.879** (-2.47, -0.91)	-1.921** (-2.61, -1.41)	-2.250** (-2.92, -1.8)
$h_{t_0} = P$	1.463 (-0.51, 3.5)	2.072** (0.17, 4.63)	2.410** (0.24, 5.44)	2.386* (-0.29, 3.9)	1.022 (-1.26, 3.32)
2 nd tercile of γ	-0.247 (-1.25, 0.3)	-0.337 (-1.26, 0.19)	-0.509** (-1.16, -0.07)	-0.546** (-0.97, -0.16)	-0.642** (-1.05, -0.28)
3 rd tercile of γ	-1.203** (-2.3, -0.54)	-1.374** (-2.55, -0.74)	-1.188** (-2.2, -0.76)	-1.286** (-1.9, -0.83)	-1.355** (-1.85, -0.91)
2 nd quintile of k_{t_0}	-0.002 (-1.12, 0.64)	-0.129 (-1.32, 0.39)	-0.048 (-0.71, 0.33)	-0.459** (-1.02, -0.15)	-0.469** (-0.99, -0.18)
3 rd quintile of k_{t_0}	-0.620 (-2.2, 0.08)	-0.429 (-1.82, 0.06)	-0.367** (-1.18, -0.07)	-0.378** (-1.06, -0.16)	-0.603** (-1.13, -0.32)
4 th quintile of k_{t_0}	-0.749* (-3.0, 0.002)	-0.606* (-2.1, 0.02)	-0.691** (-1.57, -0.23)	-0.701** (-1.23, -0.26)	-0.759** (-1.21, -0.33)
5 th quintile of k_{t_0}	-2.348** (-5.58, -1.48)	-1.616** (-4.3, -0.75)	-1.169** (-2.7, -0.8)	-1.280** (-2.36, -0.86)	-1.264** (-2.05, -0.87)

Table B2: Estimation results for the equation predicting health type, that is Eq. (6). The columns refer to specifications controlling for different number of lags of past health. We set fair health status, 1st tercile of γ and 1st quintile of k_{t_0} as the base for the corresponding dummy variables. All estimations include dummy variables for 10-years windows of birth year. The numbers in parentheses are the 95th confidence interval from bootstrapped samples. The asterisks ** (*) refer to estimates for which the 95% (90%) confidence interval excludes zero.

B.5 The distribution of productivity fixed effects

To evaluate whether the assumption of the normal distribution in our quantitative model is consistent with the data, we plot the standardized estimates of the fixed labor productivity in Figure B2, together with the fitted curve of a standard normal distribution. The graph shows that the distribution of our estimated γ is reasonably well approximated by the fitted normal distribution.




Figure B2: Distribution of estimated fixed productivity.

B.6 Estimating the two-year health process (younger starting ages)

Tables B3 and B4 report the estimation results from Eq. (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 B3: Estimation results for the health process in Eq.(2)-(3) in the top panel and Eq.(4)-(5) in the bottom panel when the sample includes only individuals whose age t_0 is less or equal to 39 years old. The columns refer to specifications controlling for different number of lags of past health. The terms $a_{\eta_2}^B$ and $a_{\eta_2}^G$ are normalized to zero. Being in bad/good health for one period ($\tau_B = 1, \tau_G = 1$) is 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	-0.162*	-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 B4: Estimation results to predict health type in Eq.(6) when the sample includes only individuals whose age t_0 is less or equal to 39 years old. The columns refer to specifications controlling for different number of lags of past health. Fair health status, 1st tercile of γ and 1st quintile of k_{t_0} are the base for the corresponding dummy variables. All estimations include dummy variables for 10-years windows for birth year. (* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$)

B.7 Estimating the two-year health 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 B5 and B6.

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 B6 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 B6 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 B5: Estimation results for the health process in Eq.(2)-(3) in the top panel and Eq.(4)-(5) in the bottom panel when the sample includes all education groups. The columns refer to specifications controlling for different number of lags of past health. The terms $a_{\eta_2}^B$ and $a_{\eta_2}^G$ are normalized to zero. Being in bad/good health for one period ($\tau_B = 1, \tau_G = 1$) is 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	-0.140*	-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 B6: Estimation results to predict health type in Eq.(6) when the sample includes all education groups (less than high school, high school, and college). The columns refer to specifications controlling for different number of lags of past health. Less than high school, fair health status, 1st tercile of γ and 1st quintile of k_{t_0} are the base case for the corresponding dummy variables. All estimations include dummy variables for 10-years windows of birth year. (* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$)

B.8 The two-year health transition probability ($T = 3$)

Table B7 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
$Pr(h_{21} = F, \tau_{21} = 1 \gamma)$	0.084	0.062	0.026
$Pr(h_{21} = G, \tau_{21} = 3 \gamma)$	0.926	0.938	0.973

Table B7: 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 B3 reports our estimated transition probabilities for all health states when $T=3$. It features very large differences by health types. Figure B4 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.

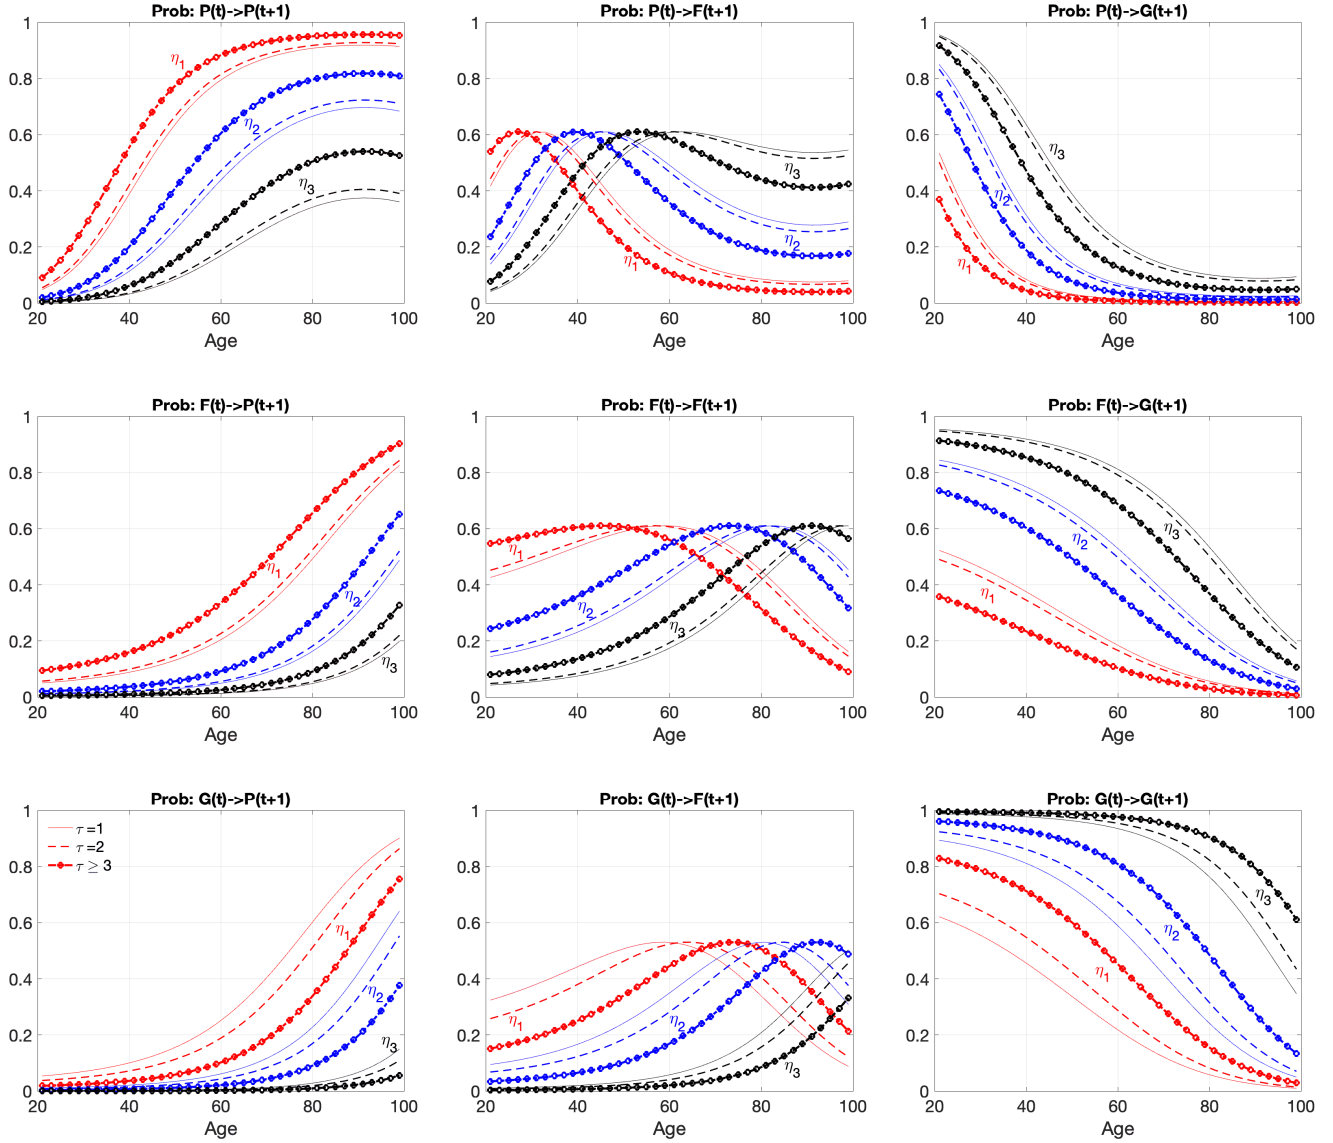


Figure B3: Estimated two-year transition probabilities ($T = 3$). On the left hand side, the top three lines refer to the low health types (η_1), the middle three lines to the middle health type (η_2) and the bottom three lines to the best health types (η_3). On the right hand side, the ordering of these groupings is reversed. Within each grouping, different line types correspond to the number of consecutive periods in current health status (τ), with the solid line referring to $\tau = 1$, the dashed line to $\tau = 2$, and the line with markers to $\tau = 3$. 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 B4: Fraction of people in bad health

B.9 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 B5 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 panel of Figure B5 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 B5: Distribution by lifetime unhealthy periods (age 21-80). Left panel: distribution among all individuals. Right panel: cumulative distribution of individuals with different health 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 25th, 50th, and 75th 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 25th 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
<i>25th percentile</i>				
0-1	-0.761	-0.692	-0.665	-0.761
4-5	-1.610	-0.487	-0.725	-0.950
<i>50th percentile</i>				
0-1	-0.120	0.036	0.003	-0.028
4-5	-0.682	0.127	0.215	-0.332
<i>75th percentile</i>				
0-1	0.473	0.720	0.666	0.560
4-5	-0.075	0.734	1.089	0.486

Table C8: The 25th, 50th, and 75th 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 and Heckman (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. Several 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 ex-ante heterogeneity is an important determinant of health dynamics.

C.3 Health types determination when including exercise behavior in health transitions

Table C9 reports the estimates of one’s health type prediction when a dummy variable for exercise categories is included in the evolution of the health process (see Section 3).

	T=5	T=4	T=3	T=2	T=1
age t_0	-0.0418	-0.0315	-0.0371	-0.0334	-0.0395
$h_{t_0} = G$	-1.886***	-2.282***	-3.001***	-2.841***	-2.943***
$h_{t_0} = P$	0.948	1.349	1.726	1.996	-0.448
2 nd tercile of γ	-0.223	-0.309	-0.440*	-0.509**	-0.543**
3 rd tercile of γ	-1.056**	-1.085***	-1.097***	-1.196***	-1.215***
2 nd quintile of k_{t_0}	0.0377	-0.122	0.033	-0.277	-0.281
3 rd quintile of k_{t_0}	-0.557	-0.177	-0.083	-0.136	-0.294
4 th quintile of k_{t_0}	-0.441	-0.308	0.015	-0.077	-0.243
5 th quintile of k_{t_0}	-2.092***	-0.994**	-1.194***	-1.282***	-1.383***

Table C9: Estimation results for one’s health type prediction when allowing for exercise to affect health. Fair health status, 1st tercile of γ and 1st quintile of k_{t_0} are the base case for the corresponding dummy variables. All estimations also include dummy variables for 10-years windows of birth year. (* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$)

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 D6 shows the medical costs for each grid separately for each health status.




Figure D6: Two-year Medical expense grids by health status, x_t^h .




Figure D7: Private health insurance coverages: $cvg(x_t^h, i_H)$, $i_H \in \{1, 2\}$.

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 D7 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).³

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

³ 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.

set the maximum (annual) taxable income for Social Security (\bar{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 \bar{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.⁴

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).⁵

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 empirical autocovariance matrix and the corresponding matrix implied by Eq.(13).⁶ Our resulting estimates of annual labor income shock are $\rho_\nu = 0.947$, $\sigma_\varepsilon^2 = 0.02$, $\sigma_{\nu_0}^2 = 0.09$, and $\sigma_\gamma^2 = 0.051$, 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 D10 below summarizes parameters of our life-cycle model.

⁴ 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.

⁵ Note that the parameters in Eq. (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.

⁶ This is a standard procedure commonly used in the literature. See for example, Storesletten et al. (2004) and French (2005).

Parameter name	Notation	Value	Source
<u>Parameters set outside the model</u>			
Risk aversion	ρ	3.0	
Average family size	\bar{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_ε^2	0.0198	"
- Initial distribution	$\sigma_{\nu_0}^2$	0.093	"
- Fixed effects	σ_γ^2	0.051	"
Proportional load in ind ins	ξ	1.07	
Fixed load in ind ins (per year)	φ^h		
- healthy		\$100	Pashchenko and Porapakarm (2017)
- unhealthy (poor+fair)		\$2,100	"
<u>Parameters used to match some targets</u>			
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)	\bar{c}	\$3,505	"
Per-period utility of being alive	\bar{b}	7.149	VSL of \$2M

Table D10: Parameters of the baseline model

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 (2007b), 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 losses of bad health when the interest rate is zero

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 E11 and E12. 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 E15. The overall patterns are similar to the case when using 2% interest rate.

	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 (% of average earnings) ^a	\$3,303 (8.5%)	\$16,800 (43%)	\$4,613 (12%)	\$627 (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 (% of average earnings) ^a	\$2,026 (5.2%)	\$13,267 (34%)	\$2,569 (6.6%)	\$130 (0.3%)

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

Table E11: 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 E12: Composition of annual monetary loss due to bad health (poor+fair) using 0% interest rate when computing the present value.

E.3 Welfare losses by type

	all	β_L			β_H		
		η_1	η_2	η_3	η_1	η_2	η_3
γ_L	\$2,113 (13.2%)	\$5,702 (35.0%)	\$2,083 (12.5%)	\$202 (1.2%)	\$6,903 (47.6%)	\$3,881 (25.2%)	\$946 (6.0%)
γ_M	\$2,003 (10.7%)	\$6,314 (33.5%)	\$2,320 (12.3%)	\$248 (1.3%)	\$8,295 (47.8%)	\$4,614 (24.9%)	\$1,202 (6.3%)
γ_H	\$1,684 (8.0%)	\$6,232 (29.6%)	\$2,341 (11.4%)	\$246 (1.2%)	9,774 (48.7%)	\$5,370 (25.1%)	\$1,433 (6.6%)

Table E13: Welfare losses due to bad health (poor+fair). The dollar value is calculated from $\lambda_c \bar{c}^{**}$ where λ_c and \bar{c}^{**} are the percentage of consumption reduction and average life time consumption of each individual when always healthy. The percentage of consumption compensation (λ_c) is reported in parentheses.

Table E13 displays the welfare losses by patience, health, and productivity type.

E.4 Alternative calibration with VSL of \$6 millions

In our baseline parameterization, we adjust the scaling parameter \bar{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 \bar{b} is set to match a VSL of \$6 millions. Note that all other parameters in the model are the same as in the baseline. The welfare losses when the targeted VSL is set to a higher value are reported in Table E14. 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 E15 reports the concentration of welfare losses and its variation due to health types. A comparison with Table 17 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 (% consumption equivalence, λ_c)	\$2,320 (12.6%)	\$6,416 (37.2%)	\$3,025 (16.6%)	\$1,321 (6.8%)	\$1,860 (10.3%)	\$2,896 (15.5%)
<u>Contribution (%)</u>						
- 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 E14: Welfare loss due to bad health (poor+fair) using VSL=\$6M. The dollar value is calculated from $\lambda_c \bar{c}^{**}$ where λ_c and \bar{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%	
<u>Monetary losses (21-death) using 0% interest rate</u>				
- Income losses + total medical costs	32%	49 %	71%	69%
<u>Welfare losses using VSL=\$6M</u>				
- Compensated consumption equivalence	23%	43%	74%	17%

Table E15: 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.