University of Connecticut

Department of Economics Working Paper Series

How Important Is Health Inequality for Lifetime Earnings Inequality?

by

Roozbeh Hosseini
University of Georgia
Federal Reserve Bank of Atlanta

Karen A. Kopecky
Federal Reserve Bank of Atlanta
Emory University

Kai Zhao
University of Connecticut

Working Paper 2020-20
December 2020

365 Fairfield Way, Unit 1063
Storrs, CT 06269-1063
Phone: (860) 486-3022
Fax: (860) 486-4463
http://www.econ.uconn.edu/

How Important Is Health Inequality for Lifetime Earnings Inequality?*

Roozbeh Hosseini[†]
University of Georgia
FRB Atlanta

Karen Kopecky[‡]
FRB Atlanta
Emory University

Kai Zhao[§]
University of Connecticut

December 2020

Abstract

Using a dynamic panel approach, we provide empirical evidence that negative health shocks reduce earnings. The effect is primarily driven by the participation margin and is concentrated in less educated and poor health individuals. We build a dynamic, general equilibrium, lifecycle model that is consistent with these findings. In the model, individuals, whose health is risky and heterogeneous, choose to either work, or not work and apply for social security disability insurance (SSDI). Health impacts individuals' productivity, SSDI access, disutility from work, mortality, and medical expenses. Calibrating the model to the United States, we find that health inequality is an important source of lifetime earnings inequality: nearly 29 percent of the variation in lifetime earnings at age 65 is due to the fact that Americans face risky and heterogeneous lifecycle health profiles. A decomposition exercise reveals that the primary reason why individuals in the United States in poor health have low lifetime earnings is because they have a high probability of obtaining SSDI benefits. In other words, the SSDI program is an important contributor to lifetime earnings inequality. Despite this, we show that it is ex ante welfare improving and, if anything, should be expanded.

Keywords: earnings, health, frailty, inequality, disability, dynamic panel estimation, life-cycle models

JEL Classification numbers: D52, D91, E21, H53, I13, I18

*We thank Jordan Herring for outstanding research assistance. We are grateful to Hannes Schwandt and David Wiczer for insightful discussions. We thank R. Anton Braun, Mariacristina De Nardi, Eric French, Dirk Krueger, Kjetil Storesletten, and Richard Rogerson for feedback. We also thank attendees at the IFS conference on Inequality and Transfers over the Life Cycle, FRB San Francisco conference on Micro-Macro Labor Economics, Barcelona GSE Summer Forum, MRDRC Researcher Workshop, Shanghai Macro Workshop, SED, Midwest Macro, and SEA meetings, as well as, seminar participants at Bank of Canada, Purdue, PHBS, UAB-IAE, UCSB, University of Hawaii, University of Minnesota, McMaster University, USC Marshall School, University of Alabama, FRB Philadelphia, Vanderbilt, William and Mary, University of Houston, and University of Tokyo. See the Online Appendix for supplemental material.

[†]email: roozbeh@uga.edu

[‡]email: karen.kopecky@atl.frb.org

[§]email: kai.zhao@uconn.edu

1 Introduction

In this paper we quantify the impact of health inequality on lifetime earnings inequality. In particular, we assess how much of the variation in lifetime earnings among older individuals in the United States (U.S.) is due to the fact that they face heterogeneous and risky life cycle health profiles.

We also assess the relative contributions of five channels through which health may impact individuals. First, poor health is associated with higher out-of-pocket medical expenses. Second, it is associated with higher mortality risk. Third, individuals in poor health may have a higher cost of working, both physically and psychologically. Fourth, poor health may have direct negative impacts on labor productivity and wages. Finally, in the U.S., individuals who are not working and in poor health have a higher likelihood of being awarded Social Security Disability Insurance (SSDI) benefits.

To quantify the impact of health inequality on lifetime earnings inequality, we build and parametrize a heterogeneous agent life cycle model. Each of the five channels through which health may impact individuals is present in the model. We find that health inequality has a large effect on lifetime earnings inequality. If we give all individuals in the economy the average life cycle health profile, the variance of log lifetime earnings at age 65 declines by 29 percent. Through a series of decomposition exercises, we find that the results are mainly driven by a combination of the SSDI, productivity, and disutility channels. The fact that individuals in poor health can obtain SSDI benefits has the largest impact, while the negative impact of poor health on labor productivity has the second largest.

Having found that the SSDI program is the primary channel through which health inequality generates lifetime earnings inequality, we then ask whether individuals in our model economy would be better off without it. We find that, even though the SSDI program increases inequality in earnings and consumption, it is welfare improving. Together with the tax implications of re-balancing the government budget, long-run ex ante welfare falls by 0.84 percent if the program is removed. The negative welfare effects are due to welfare losses of less educated individuals, while college graduates are slightly better off. Finally, we show that, if anything, individuals would prefer if the SSDI program were more generous. Raising SSDI benefit levels by 10 percent in the model increases ex ante welfare by 0.20 percent.

Our analysis employs a new objective measure of health status, called the *frailty index*.[1] The frailty index is simply the accumulated sum of all adverse health events that an individual has incurred. Each health problem is referred to as a *deficit*. An important advantage of the frailty index is that it measures health on a fine scale. There is significant variation in frailty among working-age Americans. But, as we show in Hosseini et al. (2019), most of this variation is concentrated among unhealthy individuals. In other words, the cross-sectional distribution of frailty is highly-skewed. We document that, at any age, moving up along the thin unhealthy tail of the frailty distribution is correlated with lower probabilities of being employed and a higher probability of being on SSDI.

To start, we conduct an empirical analysis that is used to motivate and guide the development of our structural model. In particular, we use data from the Panel Study of Income

[1] We are not the first study that uses the frailty index as a measure of health status. See Dalgaard and Strulik (2014), Schünemann et al. (2017b), and Schünemann et al. (2017b,a) for other instances.

Dynamics (PSID) and a dynamic panel data approach (see Blundell and Bond (1998)) to estimate the impact of health on current earnings and its components: participation, hours conditional on working, and wages. The fact that we can measure frailty on a finer scale also means that we can treat it as a continuous variable. This is an attractive feature of the frailty index for our empirical analysis because it allows us to use it to estimate the marginal impacts of health on earnings and its components.

In our dynamic panel data estimation, we find that an incremental deterioration of health (adding one more deficit to frailty) reduces earnings by 20 percent. The impact operates primarily through the employment margin and is more pronounced among individuals who are already in poor health and those without a college degree. In particular, we find no statistically significant effect on hours conditional on working. The effects of frailty on wages are smaller and only significant for individuals without a college degree. We also use our dynamic panel estimator to estimate the causal effect of changes in earnings on frailty. Controlling for age and fixed effects, we fail to find statistically significant effects both overall and, on average, within the education and health groups we consider.

Using our empirical findings as a guide, we build a general equilibrium, life cycle model featuring agents who experience heterogeneous and risky *frailty* dynamics over their life cycle as well as productivity and employment risk. Agents jointly make consumption, savings, and labor supply decisions in each period over their life cycle. Given our empirical findings, we assume that individuals in the model only adjust labor supply on the extensive margin. Working-age individuals can choose to work or exit the labor force and apply for disability insurance. Retirement-age individuals can choose to work or retire. In the model, markets are incomplete, but there exists a government that runs the disability insurance program, as well as, a social security program, and a tax/transfer system.

Since we do not find any statistically significant effects of earnings on frailty, we do not allow for such feedback effects in the model. Instead, we assume that individuals in the model face exogenously-given frailty processes. An individual's frailty affects their behavior through five different channels: 1) mortality rates, 2) out-of-pocket medical expenditures, 3) labor productivity, 4) probability of successful DI application, and 5) disutility from working. We estimate the effect of frailty on the first three directly from the data. To estimate the effect of frailty on productivity we use our dynamic panel data estimator and a selection correction procedure proposed by Al-Sadoon et al. (2019).

The effects of the second two channels are pinned down using the model and a simulated method of moments procedure. To disentangle the effect of the SSDI channel from that of the disutility one we use the following strategy. We target the variation in SSDI recipiency rates and labor force participation rates by frailty for 5-year age groups of individuals ages 25 to 64. However, we also target the dispersion in labor force participation rates by frailty for individuals aged 65 to 84. The variation in labor force participation by frailty after age 65 cannot be directly due to the SSDI channel. In the U.S. after age 65 individuals can no longer receive SSDI benefits. Instead, all individuals (regardless of their health status), are eligible for social security retirement benefits. Thus, these additional moments pin down the effect of frailty on the disutility of working. Since most of the variation in labor supply and SSDI recipiency occurs in the unhealthy tail of the frailty distribution, the set of moments targeted is concentrated in this tail.

To assess the quality of our baseline calibration we use a set of non-targeted moments.

The moments we focus on are the variation in labor force participation rates and SSDI recipiency rates by education. We have three education groups in the model: high school dropouts, high school graduates, and college graduates. We compare both the aggregate labor force participation rates and SSDI recipiency rates by education in the model to the data and the variation in these rates by frailty and age within each education group. The model is able to replicate the patterns in these rates observed in the data.

We use the calibrated model to run the following counterfactual experiment. We assign the average age profile of frailty to all individuals in the model and compare the variance of log lifetime (cumulative) earnings in the counterfactual model to that in the baseline. Removing health inequality in this way reduces the variance of log lifetime earnings by 13 percent at age 45, 27 percent at age 55, 29 percent at age 65, and 21 percent at age 75. Inspection of the ratios of lifetime earnings at the 5th, 10th, 90th and 95th percentiles relative to the median reveals that the impact of health inequality on lifetime earnings inequality is concentrated in the bottom of the income distribution.

Then we conduct a series of counterfactual exercises to explore the relative importance of the various channels through which health operates in the model. In each experiment we turn off the effect of frailty only in one channel. We find that, at younger ages, the decline in lifetime earnings inequality when we remove health inequality is primarily due to the labor productivity channel. However, at older ages, it is primarily, and largely, due to the DI channel. Relative to the labor productivity and DI channels, the other three channels have a relatively small impact. In particular, the effect of frailty on disutility of work does not seem to be an important determinant of how health affects labor supply and earnings inequality.

The reason why the primary channel through which health inequality operates differs by age is as follows. In the baseline economy, poor health negatively impacts the productivity of less-educated workers only. As a result, using the average frailty profile to determine individuals' labor productivity reduces lifetime earnings inequality at all ages. In contrast, assuming that DI eligibility is determined by the average frailty profile, has two effects. One, it reduces the incentive for frail young individuals to work. They no longer have a high probability of getting on DI when older. Thus, instead of accumulating lifetime earnings to increase their expected future DI benefits, they prefer to avoid the disutility of work and rely on means-tested welfare programs. Two, it increases the labor supply of older frail individuals. These individuals no longer have a high probability of getting DI either, but, because they have more accumulated wealth, they are less likely to obtain means-tested benefits if they stop working. Thus, they prefer to work until retirement age. The importance of the second effect grows larger with age making the DI channel the most important channel through which health inequality increases lifetime earnings inequality from age 55 on.

Finally, we explore the welfare implications of changes to the SSDI program. Removing the SSDI program reduces ex ante welfare, despite increasing aggregate output and consumption and reducing earnings inequality. The ex ante welfare loss is due to highly frail less educated workers who have low labor productivity and high disutility from work. Without the SSDI program, they choose between increased dependence on means-tested transfers for consumption or working more despite the high utility costs and relatively low returns.

Our paper belongs to the growing literature that uses rich life cycle models to study the aggregate and distributional economic impacts of health status and health expenditures.[2] Two closely related papers are Capatina (2015) and Low and Pistaferri (2015). Capatina (2015) studies the effect of poor health on labor supply using a life cycle framework that is similar to ours. The most important difference between our study and hers is that we explicitly model SSDI and the incentive effects it has on labor supply. This allows us to disentangle the increased disutility of working effect of declines in health from the increased SSDI access effect. We find that increased access to SSDI is the most important channel through which health inequality impacts lifetime earnings inequality. Unlike Capatina (2015), we find that the effect of health on preferences only plays a small role. Low and Pistaferri (2015) develop a life cycle model similar to ours but for a different purpose. While we focus on health and earnings inequality, they study the impact of reforms of SSDI and means-test welfare programs on welfare and behavior. Our findings on the welfare implications of cutting/raising SSDI benefits by 10 percent are consistent with their analysis.

Other related papers which focus on the SSDI program include Kitao (2014) who studies the joint effect of SSDI benefits and Medicare eligibility on SSDI recipiency and labor supply and Michaud and Wiczer (2017) who use a model to measure the impact of health deterioration and concentration of health risks within certain occupations on SSDI claims. Also, Kim and Rhee (2020) who study the impact of temporarily restricting access to SSDI in a model where agents have heterogeneous human capital and Aizawa et al. (2020) who derive an optimal disability policy in a model where firms strategically vary access to work flexibility in order to screen out disabled workers.

Our paper is also related to the empirical literature that estimates the effect of health and SSDI recipiency on employment and earnings. For instance, Blundell et al. (2017) find that 15 percent of the decline in employment from age 50 to 70 is due to declines in health in the US and UK. Similar to us, they find the effect to be larger among less educated workers. Bound et al. (1999) show that deteriorating health is associated with declines in labor force participation using the Health and Retirement Study. Meyer and Mok (2019) find that the prevalence of disability in the U.S. is high and correlated with poor economic outcomes including lower earnings and labor supply. Autor and Duggan (2003) document that the labor force participation rates of high school dropouts are responsive to the stringency and generosity of the SSDI program and Maestas et al. (2013) show that the labor force participation rates of marginal SSDI applicants are sensitive to successful benefit receipt.[3]

The remainder of the paper is organized as follows. In Section 2 we document empirical facts on the relationship between health status and earnings. These facts are used to guide the development of the model we present in Section 3. The calibration of the model is outlined in Section 4. In Section 5 we assess the model's ability to replicate non-targeted moments. Section 6 reports the results of our quantitative exercise and Section 7 concludes.

[2]See French (2005), De Nardi et al. (2010), De Nardi et al. (2017), Suen (2006), Kopecky and Koreshkova (2014), Zhao (2014), Ozcan (2013), Braun et al. (2015), Imrohoroglu and Zhao (2018), and Prados (2017), among others.

[3]See also García-Gómez et al. (2013), Lundborg et al. (2015), Dobkin et al. (2018), Heinesen and Kolodziejczyk (2013), Jeon (2017) and Pohl et al. (2013) who use specific health events to estimate the effect of health on labor market outcomes.

2 Empirical Facts on Health and Earnings

We start by documenting some empirical facts on the relationship between *health status* and earnings that we use to guide the development of our structural model. However, first, we need to introduce and motivate our measure of health status: the *frailty index*. This brief overview draws heavily on Hosseini et al. (2019) which includes additional details on the properties of the index and provides an extensive comparison between it and other commonly-used measures of health.

2.1 Frailty index as a measure of health status

As individuals age they develop an increasing number of health problems, functional impairments, and abnormalities. Some of these conditions are rather mild (e.g., reduced vision) while others are serious (e.g., heart disease). However, as the number of these conditions rises, the person's body becomes more frail and vulnerable to adverse outcomes. We refer to each of these conditions as a *deficit*. In their pioneering work, Mitnitski et al. (2001) and Mitnitski et al. (2002) demonstrated that the health status of an individual can be represented by an index variable, called the *frailty index*, which summarizes the individual's accumulated deficits. The index is constructed as the ratio of deficits a person has accumulated to the total number of deficits considered. For example, if 30 deficits were considered and 3 were present for a person, that person is assigned a frailty index of 0.1.[4]

We use three datasets to quantify the impact of health inequality on lifetime earnings inequality: the Panel Study of Income Dynamics (PSID), Health and Retirement Study (HRS) and Medical Expenditure Panel Survey (MEPS). To construct frailty indices for individuals in each dataset we use health deficit variables from three broad categories: restrictions and difficulties in Activities of Daily Living (ADL) and Instrumental ADL (IADL); mental and cognitive impairments; and medical diagnosis and measurements. Examples of deficits in the first category are difficulty eating, dressing, or walking across a room without assistance. Examples from the second category are cognitive and memory test scores from tests of abilities such as backwards counting and immediate word recall. Examples from the third category are having received a diagnoses of high blood pressure, diabetes, or obesity.

We use the frailty index to measure health for our analysis because it has several attractive properties for studying individuals' life-cycle health dynamics and their implications. First, despite its simplicity, it is well documented in Gerontology that the index is highly predictive of health outcomes. Mitnitski et al. (2004) (among others) have found that having a higher frailty index is associated with a higher likelihood of an adverse health outcome, such as death or institutionalization.[5]

[4] We show in Hosseini et al. (2019) that the properties of the index are robust to principal component weighting. One could come up with many other alternative weighting schemes. However, equal weighting is simple and works well. This may be, at least in part, because individuals with more severe conditions are likely to have more total deficits. For instance, consider two individuals with cancer. The one with a more serious case will likely also report limitations with ADL's and IADL's.

[5] See also Searle et al. (2008); Rockwood and Mitnitski (2007), Rockwood et al. (2007), Mitnitski et al. (2001), Mitnitski et al. (2005), Kulminski et al. (2007a), Kulminski et al. (2007b), Goggins et al. (2005), and Woo et al. (2005).

Table 1: Frailty Summary Statistics in our PSID Sample

Mean	0.11	Median	0.07
by age:		Standard Deviation	0.11
25-34	0.07		
35-44	0.09	$+\Delta$ Frailty	0.29
45-54	0.11	$-\Delta$ Frailty	0.13
55-64	0.15	Effect of 1 additional deficit	+0.037

Second, the frailty index measures health at a fine enough level that it can be treated as a continuous variable. This is a desirable feature for two reasons. One, it allows us to quantify the impact of marginal changes in frailty on economic outcomes as we do in our empirical analysis below. Two, the distribution of frailty is significantly right-skewed. In others words, there is substantial variation in the extent of poor health among individuals in the unhealthy right-tail. As we will show below, the impacts of declines in health on earnings and labor supply are concentrated in this unhealthy tail of the distribution. Targeting this variation when calibrating our structural model is important for accurately quantifying the impact of health on lifetime earnings inequality but doing so requires having a sufficiently fine measure of health.

2.2 Empirical strategy to estimate health effects

In this section we document some empirical facts on the relationship between frailty and earnings that are used to develop our structural model. To this end, we estimate the effect of frailty on current earnings and assess the relative importance of the three margins through which the effect may operate: labor force participation, hours conditional on working, and wages.

The sample we construct to conduct our analysis is based on the eight waves of the PSID covering the period 2002–2016.[6] Our extended PSID sample consists of household heads and spouses ages 25 to 94. However, we restrict the sample to ages 25 to 64 for the dynamic panel analysis. Table 1 provides summary statistics on frailty for individuals in the dynamic panel sample.[7] Notice that the cross-sectional distribution of frailty is right-skewed and mean frailty increases with age. We use 27 deficit variables to construct the frailty index in PSID. Thus, incurring one additional deficit increases one's frailty index by 1/27 or 3.7 percent. Wave-to-wave changes in frailty occur for 42 percent of the sample on average, 69 percent of which are increases.

Figure 1 shows the raw correlations of frailty with earnings, participation, hours conditional on working, and wages by 5-year age groups for individuals aged 25 to 74.[8] The

[6] We start our sample in 2002 because the PSID did not collect enough information on individuals' medical conditions, ADL's, and IADL's prior to the 2003 wave to construct frailty indices. The PSID is biennial over this period.

[7] Additional details on sample selection and summary statistics can be found in Section 1 of the Online Appendix.

[8] Earnings are annual labor earnings where labor earnings of non-workers are set to zero. Individuals are counted as participating in the labor force if they worked at least 260 hours during the year at a wage of at

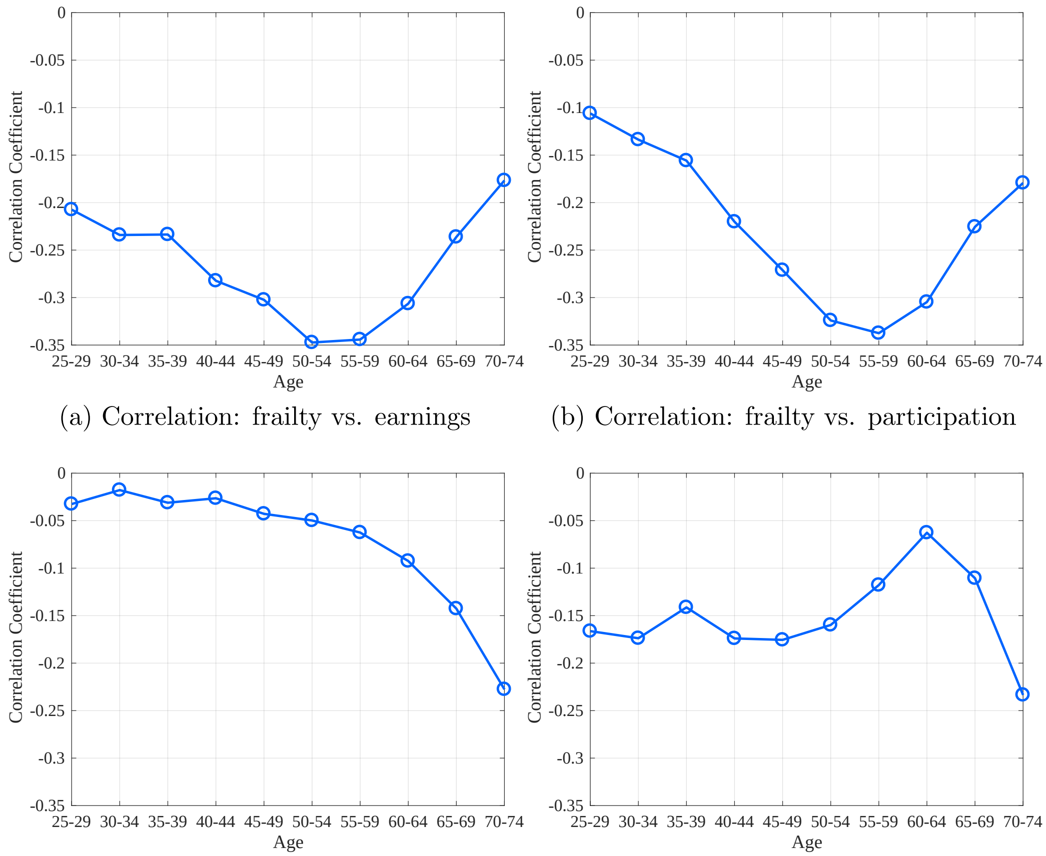

(a) Correlation: frailty vs. earnings

(b) Correlation: frailty vs. participation

(c) Correlation: frailty vs. hours (workers)

(d) Correlation: frailty vs. wages (workers)

Figure 1: Raw correlations of earnings (top left), participation (top right), hours conditional on working (bottom left), and observed wages (bottom right) with frailty for 25 to 74 year-olds by five-year age groups. Data source: 2003–2017 PSID.

figure shows that frailty and earnings are negatively correlated at all ages. The negative correlation is u-shaped over the life cycle with the magnitude peaking during the fifties. Notice that this negative correlation is due to a negative correlation between frailty and all three components of earnings (participation, hours, and wages). However, of the three, the correlation between frailty and participation is the largest and follows a similar life cycle pattern to that of earnings.

How much of the negative correlation between frailty and earnings is driven by declines in health generating declines in earnings and how important are the various margins? To answer this question, we now use a dynamic panel approach to estimate the causal impacts of frailty on earnings, hours and wages. We estimate the following statistical model

$$y_{i,t} \quad = \quad b_i + \gamma f_{i,t} + \alpha_1 y_{i,t-1} + \alpha_2 y_{i,t-2} + \delta \boldsymbol{Z}_{i,t} + \varepsilon_{i,t}, \tag{1}$$

in which $y_{i,t}$ is the logarithm of either earnings, hours, or wages for individual i at time t. $f_{i,t}$ is frailty and $\boldsymbol{Z}_{i,t}$ is a vector of exogenous controls that includes marital status, marital

least \$3 per an hour. Annual hours worked are calculated as weekly hours times weeks worked. Wages are constructed by PSID using annual labor earnings and annual hours worked.

status interacted with gender, number of kids, number of kids interacted with gender, year dummies, and a quadratic in age.[9] Finally, b_i is the individual fixed effect and $\varepsilon_{i,t}$ is a random error term with

$$E[b_i] = E[\varepsilon_{i,t}] = E[b_i\varepsilon_{i,t}] = 0.$$

Individuals vary in unobservable ways (such as innate ability) that could potentially be correlated with both their earning ability and frailty. This motivates the inclusion of a fixed effect in equation (1).[10]

We are interested in estimating the impact of frailty on earnings. However, earnings may also impact frailty. Declines in health may affect productivity and lead to lower wages (or loss of employment). Yet, lower income (or loss of employment) may negatively impact health through its impact on mental health, access to health insurance, or choice of medical care. Moreover, both earnings and frailty are highly persistent variables. In other words, we are concerned about simultaneity but also dynamic endogeneity: past earnings are correlated with current earnings but may also be correlated with both past and current frailty. This concern is the reason why we use a dynamic panel data approach. We need to include lagged values of earnings on the right-hand-side in equation (1).

It is well known that equation (1) cannot be consistently estimated using OLS or fixed effect estimators (see Nickell (1981) and Wooldridge (2010) for details). Therefore, to obtain a consistent and unbiased estimate of the effect of frailty on earnings we use a dynamic GMM panel estimator. This class of estimators was introduced by Holtz-Eakin et al. (1988) and Arellano and Bond (1991), and further developed by Blundell and Bond (1998) (and many others).[11]

The basic estimation procedure consists of two steps. The first step is to write equation (1) in first difference form:

$$\Delta y_{i,t} \;\; = \;\; \gamma\Delta f_{i,t} + \alpha_1\Delta y_{i,t-1} + \alpha_2\Delta y_{i,t-2} + \delta\Delta \boldsymbol{Z}_{i.t} + \Delta\varepsilon_{i,t} \tag{2}$$

which eliminates time-invariant unobserved heterogeneity. The second step is to use lagged values of the left-hand-side variable, frailty, and the endogenous controls in levels as 'internal' instruments and estimate equation (2) using GMM. As we argued above, lagged values of frailty and earnings are predictors of current levels of earnings and frailty. Therefore, they provide sources of variations for current values. However, for instruments to be valid, the past levels of earnings and frailty must be uncorrelated with $\varepsilon_{i,t}$. In other words, the following

[9] We have also experimented with making marital status, marital status interacted with gender, number of kids, and number of kids interacted with gender endogenous and found that it does not have a significant impact on any of our results.

[10] The individual fixed effect controls for any non-time-varying heterogeneity across individuals including, for example, differences in gender and education. While we have done are best to include all relevant controls, we cannot completely rule out the possibility that we have omitted time-varying variables that are correlated with both frailty and are left-hand-side variables.

[11] Dynamic panel estimators are widely used in many areas of economics and finance. Examples include the effect of board structure on firm performance (Wintoki et al. (2012)), capital accumulation and firm investment (Whited (1991)), the sensitivity of firm investments to available internal funds (Bond and Meghir (1994)), economic growth convergence (Caselli et al. (1996)), estimation of a labor demand model (Blundell and Bond (1998)), the relation between financial intermediary development and economic growth (Beck et al. (2000)), and the diversification discount (Hoechle et al. (2012)), among many others.

orthogonality or moment conditions must hold

$$E\left(y_{i,t-s}\Delta\varepsilon_{i,t}\right)=E\left(f_{i,t-s+2}\Delta\varepsilon_{i,t}\right)=0,\qquad \text{for } \forall s>3. \tag{3}$$

Using the moment conditions in equation (3) we can estimate equation (2) via GMM. However, there are still a few shortcomings. For example, differencing can reduce variation in explanatory variables and therefore reduce accuracy of estimates (see Beck et al. (2000)). Moreover, as Arellano and Bover (1995) point out, variables in levels may be weak instruments for first-differences. This is especially true for highly persistent variables.[12] To mitigate these shortcomings, we follow Blundell and Bond (1998) and Blundell and Bond (2000) and improve the GMM estimator by jointly estimating the equation in levels and the equation in first-differences. Lagged first-differences are used to instrument levels. More precisely, we stack levels and first differences in the following equation

$$\begin{bmatrix} y_{i,t} \\ \Delta y_{i,t} \end{bmatrix} = \gamma \begin{bmatrix} f_{i,t} \\ \Delta f_{i,t} \end{bmatrix} + \alpha_1 \begin{bmatrix} y_{i,t-1} \\ \Delta y_{i,t-1} \end{bmatrix} + \alpha_2 \begin{bmatrix} y_{i,t-2} \\ \Delta y_{i,t-2} \end{bmatrix} + \delta \begin{bmatrix} \boldsymbol{Z}_{i.t} \\ \Delta \boldsymbol{Z}_{i.t} \end{bmatrix} + \varepsilon_{i,t}, \tag{4}$$

which we can estimate using the "system" GMM estimator. Note, however, that the estimation drops the fixed effect from the levels equation. As a result, for first differences to be valid instruments of the levels, the following additional orthogonality conditions must hold

$$E\left(\Delta y_{i,t-s}\left(b_i+\varepsilon_{i,t}\right)\right)=E\left(\Delta f_{i,t-s+2}\left(b_i+\varepsilon_{i,t}\right)\right)=0,\qquad \text{for } \forall s>3. \tag{5}$$

To summarize, we carry out GMM panel estimation using the orthogonality conditions (3) and (5). These conditions imply that we can use lagged levels of our endogenous regressors ($y_{i,t}$ and $f_{i,t}$) as instruments for our differenced equations and lagged differences as instruments for the levels equations, respectively. Given the concerns about instrument proliferation and overfitting discussed in Roodman (2009) we do not use all the available lags as instruments. Instead, we use only the fourth and fifth lags for the regressions that include everyone in the sample. To obtain valid instruments for the regressions that are run only on workers requires us to go back further in lag length. Thus, for these regressions we use the fifth and sixth lags. In addition, in all regressions run, we restrict the coefficients on the lags to be the same at each time t by collapsing the instrument matrix.[13]

Following the recommendations in the literature by Roodman (2009), Bond (2002) and others, we conduct several tests of our specification, approach, and instrument set. In the tables that follow, we report test statistics for two sets of tests.[14] First, we report the results of the tests for first and second-order serial correlation in the residuals of the difference equation. By construction, the residuals of the first-differenced equation should possess first-order serial correction. However, if the assumption of serial independence in the errors in the

[12] As an stark example, imagine a random walk process. In that case, past levels are uncorrelated with first differences.

[13] This increases the power of the Hansen-Sargan test for over-identification.

[14] In Section 2 of the Online Appendix we show the results of additional tests and robustness checks including robustness checks to the number of instruments used, results from difference-in-Hansen tests on subsets of instruments in the levels equation, results of instrument power tests, and a comparison of our dynamic panel GMM estimates to both estimates obtained using an OLS estimator and a within groups (FE) estimator.

levels equation is correct, the first-differenced residuals should not exhibit significant AR(2) behavior. Thus, if we pass the test for second-order serial correlation, it means that we have included enough lags to control for the dynamic aspects of our empirical relationship. As a result, any historical value of earnings beyond those lags is a potentially valid instrument since it will be exogenous to current earning shocks.

The second set of tests statistics we report are tests of the validity of our instruments. The system is over-identified in that we have more instruments than we do endogenous regressors. We conduct a Hansen-Sargan test and report the Hansen J-statistic. The Hansen J-statistic is distributed χ^2 under the null hypothesis that there is no correlation between the error terms and the instruments. Finally, we report the test statistic for the difference-in-Hansen test that our lagged first-difference instruments are uncorrelated with the fixed effects. This must hold for lagged first-differences to be valid instruments of the endogenous variables in the levels equation since the fixed effect is still in the error term.[15]

2.3 Estimation results

We now present the results from our dynamic panel estimation. We report two sets of estimation results that highlight the differences in the effect of changes in frailty on the intensive versus the extensive margin of labor supply. The first set shows the results from estimating equation (1) for everyone in our PSID sample, regardless of their labor force participation status.[16] The second set reports results only for those who are working in all periods we observe them. To aid in the interpretation of their magnitude, we report all the estimated frailty effects (and their standard errors) as the impact of the accumulation of one additional deficit. To achieve this, we rescale all the coefficients on terms involving frailty by 1/27.[17] For instance, the third row of column (1) in Table 2 reports $\gamma/27$.

Table 2 reports the results from our system GMM estimation of equation (1) where the left-hand-side variable is log earnings. Columns (1) through (4) show the regression results for the entire sample. Columns (5) through (8) show results only for workers. Notice that the p-values from the AR(2), Hansen, and difference-in-Hansen tests are all above 5 percent. Thus, we cannot reject the nulls of no second-order serial correlation in the error terms and instrument validity.

As column (1) of Table 2 indicates, frailty has a large and statistically significant effect on earnings. Accumulating one more deficit reduces earnings, on average, by 20 percent. However, as column (5) shows, the effect of accumulating one additional deficit is much smaller when conditioning on those who continue to work. For workers, the overall effect is a 4 percent decline and is borderline significant. These findings suggest that, consistent with the magnitudes of the raw correlations in Figure 1, the effects of frailty on earnings are due primarily to the extensive margin (of unhealthy workers leaving employment), rather than the intensive margin (of unhealthy workers working fewer hours or receiving lower wages).[18]

[15] There can still be correlation between the levels and the unobserved effects but this correlation must be constant over time.

[16] To take logarithms, we shift all observations of annual earnings up by \$1.00.

[17] Recall that in our PSID sample we have a total of 27 potential deficits so accumulating one more deficit is equivalent to an increase in frailty by $\frac{1}{27}$.

[18] Since exit from the labor force due to bad health is likely often permanent, we focus our discussion on

Table 2: Effect of Frailty on Earnings

	Everyone				Workers			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\log(\text{earnings}_{t-1})$	0.283	0.370	0.220	0.628**	1.474***	1.371***	1.293***	1.127***
	(0.364)	(0.319)	(0.362)	(0.291)	(0.509)	(0.400)	(0.410)	(0.302)
$\log(\text{earnings}_{t-2})$	0.396	0.318	0.444	0.115	-0.640	-0.569	-0.498	-0.308
	(0.298)	(0.259)	(0.297)	(0.239)	(0.454)	(0.356)	(0.377)	(0.273)
frailty_t	-0.199***				-0.036**			
	(0.061)				(0.017)			
$\text{frailty}_t \times$ HSD		-0.232***				-0.068**		
		(0.066)				(0.030)		
$\text{frailty}_t \times$ HS		-0.207***				-0.046***		
		(0.058)				(0.002)		
$\text{frailty}_t \times$ CL		-0.093*				-0.021		
		(0.052)				(0.018)		
$\text{frailty}_t \times$ Good Health			-0.071				-0.065	
			(0.178)				(0.066)	
$\text{frailty}_t \times$ Bad Health			-0.193***				-0.036**	
			(0.065)				(0.017)	
$\text{frailty}_t \times$ Young				-0.185***				-0.061**
				(0.066)				(0.025)
$\text{frailty}_t \times$ Old				-0.149***				-0.011
				(0.049)				(0.014)
Controls	YES	YES	YES	YES	YES	YES	YES	YES
Observations	64,965	64,965	64,965	64,965	34,274	34,274	34,274	34,274
AR(1) test (p-value)	0.455	0.319	0.497	0.104	0.030	0.010	0.021	0.008
AR(2) test (p-value)	0.380	0.474	0.298	0.949	0.130	0.082	0.138	0.160
Hansen test (p-value)	0.796	0.132	0.826	0.752	0.434	0.826	0.543	0.465
Diff-in-Hansen test (p-value)	0.652	0.360	0.827	0.464	0.255	0.484	0.259	0.214

Notes: Columns (1)–(4) show regression results for the entire sample, regardless of employment status. Columns (5)–(8) show results conditional on continued employment. All the frailty effects are reported as the effect of one additional deficit. All regressions include controls (marital status, marital status interacted with gender, number of kids, number of kids interacted with gender, year dummies, and a quadratic in age). 'HSD' is high school dropout, 'HS' is high school graduate, and 'CL' is college graduate. 'Good/Bad Health' is frailty below/above the 75th percentile. 'Young/Old' are individuals younger/older than 45 years of age. Standard errors are in parenthesis. *$p < 0.1$; **$p < 0.05$; ***$p < 0.01$.

Columns (2) and (6) of Table 2 show how the effects of frailty on earnings differ by education groups: high school dropouts (HSD), high school graduates (HSG), and college graduates (CL). Looking at column (2), frailty has a large and strongly significant effect on the earnings of high school dropouts and those with no more than a high school degree. However, the effect is smaller and less significant for college graduates. One additional deficit reduces the earnings of high school dropouts, high school graduates, and college graduates by 23, 21, and 9 percent, respectively. Among workers, the effects by education group are considerably smaller and the levels of significance are lower. For instance, accumulating one more deficit reduces the earnings of high-school dropout workers by 7 percent and the earnings of high-school graduates by 5 percent.

Columns (3) and (7) of Table 2 show results by health status. Column (3) shows that frailty has a significant and large effect on the earnings of individuals with bad health. Here,

the estimated short-run effects of frailty on earnings and not longer-run effects due to earnings persistence.

individuals with bad health are those with a value of frailty above the 75th percentile of the overall frailty distribution. For these individuals one more deficit leads to a 19 percent drop in earnings. The effect for individuals with good health (those with a value of frailty at or below the 75th percentile) is small and insignificant. Once we restrict the sample to workers, the coefficient on frailty interacted with bad health becomes substantially smaller and less significant. These results show that the negative effects of frailty on earnings are primarily due to exit from the labor force in response to additional health declines by individuals who are already in poor health.

Finally, columns (4) and (8) of Table 2 show the effects by age groups. Somewhat surprisingly, the effects of frailty on earnings are slightly smaller for those older than 45 than those younger than 45. One additional deficit reduces earnings of the young by 18 percent and earnings of the old by 15. In addition, once the sample is restricted to only those who continue to stay employed, the effect is only significant for younger workers. These estimates reveal that the effect of health on earnings for older workers operates primarily through the extensive margin. In contrast, there are potential effects operating through the intensive margin for younger workers. This suggests that, relative to older workers, younger workers may be more willing to continue working despite incurring reductions in labor earnings due to poor health.

The results in Table 2 tell us that frailty has a large and significant effect on the earnings of low educated individuals and those with bad health. They also indicate that the effect is mainly along the extensive margin. The impact on earnings of workers is either not significant or is relatively small. However, these results do not tell us whether the impacts on workers come from hours or wages, or both. To better understand how frailty affects hours and wages of workers, we repeat the same regressions but replace the independent variable with hours and/or wage. The results of these regressions are reported in Table 3.

The left panel in Table 3 (Panel A) shows the estimated effects of frailty on hours worked. The results in columns (1) to (4) show that, for the entire sample, the impact of accumulating one additional deficit on hours is very similar to the impact on earnings. Overall, accumulating one more deficit cause a 14 percent drop in hours. As with earnings, high school dropouts and those in bad health experience the largest declines in hours. College-educated and those in good health have the smallest impact.

Columns (5) through (8) of the table show the effect of frailty on hours for workers only. Note that we do not find any evidence of a significant effect either overall or in any of the subgroups we consider. This indicates that the impact of frailty on hours worked is almost entirely through exit from employment (as opposed to adjustment of working hours). In other words, if an adverse health event does not drive a worker out of employment, there will be no adjustment in hours worked, even among the high school dropouts (who experience the largest decline in earnings while working).

The right panel in Table 3 (Panel B) shows regression results for the effect of frailty on wages of workers. On average, one more deficit reduces wages of workers by 2 percent. The effect is negative for all three education groups but decreases with education and is only significant for less educated workers (those without a college degree). Of all subgroups we consider, the effect is largest for the high school dropout group. One additional deficit reduces their hourly wages by 7 percent. Notice that the effects on wages are very similar in magnitude to the effects on earnings of workers reported in Table 2.

Table 3: Effect of Frailty on Hours and Wage

	Panel A. Hours regression								**Panel B.** Wage regression			
	Everyone				Workers				Workers			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(1)	(2)	(3)	(4)
$\log(\text{hours}_{t-1})$	0.399	0.383	0.386	0.669***	0.003	0.074	0.040	0.382				
	(0.322)	(0.319)	(0.317)	(0.257)	(0.345)	(0.313)	(0.311)	(0.318)				
$\log(\text{hours}_{t-2})$	0.263	0.269	0.272	0.048	0.304	0.168	0.282	0.254				
	(0.257)	(0.253)	(0.253)	(0.206)	(0.218)	(0.221)	(0.219)	(0.246)				
frailty_t	-0.144***				0.003				-0.023**			
	(0.044)				(0.009)				(0.010)			
frailty_t × HSD		-0.177***				-0.001				-0.069***		
		(0.049)				(0.013)				(0.023)		
frailty_t × HS		-0.159***				0.001				-0.033***		
		(0.045)				(0.010)				(0.011)		
frailty_t × CL		-0.082**				0.009				-0.008		
		(0.041)				(0.009)				(0.011)		
frailty_t × Good Health			-0.082				-0.002				0.013	
			(0.041)				(0.034)				(0.062)	
frailty_t × Bad Health			-0.137***				0.001				-0.022*	
			(0.046)				(0.010)				(0.012)	
frailty_t × Young				-0.132***				-0.011				-0.041**
				(0.049)				(0.014)				(0.017)
frailty_t × Old				-0.116***				0.005				-0.015
				(0.035)				(0.010)				(0.011)
$\log(\text{wage}_{t-1})$									0.212	0.122	0.303	0.511
									(0.541)	(0.368)	(0.449)	(0.399)
$\log(\text{wage}_{t-2})$									0.532	0.600*	0.461	0.272
									(0.489)	(0.328)	(0.419)	(0.359)
Controls	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Observations	64,965	64,965	64,965	64,965	34,274	34,274	34,274	34,274	34,170	34,170	34,170	34,170
AR(1) test (p-value)	0.287	0.290	0.289	0.043	0.409	0.286	0.335	0.180	0.651	0.518	0.552	0.362
AR(2) test (p-value)	0.596	0.569	0.565	0.706	0.273	0.572	0.312	0.642	0.454	0.189	0.474	0.734
Hansen test (p-value)	0.971	0.317	0.838	0.811	0.060	0.166	0.174	0.051	0.085	0.374	0.207	0.170
Diff-in-Hansen test (p-value)	0.944	0.597	0.713	0.545	0.080	0.062	0.108	0.037	0.044	0.145	0.082	0.104

*Notes:*Panel A (left) shows regression results for the effect of frailty on hours worked. Columns (1)–(4) show regression results for the entire sample, regardless of employment status. Columns (5)–(8) show results conditional on continued employment. Panel B (right) shows regression results on the effect of frailty on wage for workers only. All the frailty effects are reported as the effect of one additional deficit. All regressions include controls (marital status, marital status interacted with gender, number of kids, number of kids interacted with gender, year dummies, and a quadratic in age). 'HSD' is high school dropout, 'HS' is high school graduate, and 'CL' is college graduate. 'Good/Bad Health' is frailty below/above the 75th percentile. 'Young/Old' are individuals younger/older than 45 years of age. Standard errors are in parenthesis. $*p < 0.1$; $**p < 0.05$; $***p < 0.01$.

Finally, as we mentioned above, it is possible that, while changes in health impact earnings, changes in earnings also impact health. In Section A.1 of the Appendix, we explore this possibility empirically. Specifically, using our dynamic panel estimation procedure, we run a similar set of regressions to those in Table 2 except with frailty on the left-hand-side and log earnings on the right-hand-side. In all cases, controlling for age and fixed effects, we find no evidence of a statistically significant effect of changes in earnings on frailty, either overall or in any of the subgroups considered.[19]

To summarize, our empirical analysis yields the following five findings. One, rises in frailty significantly reduce earnings and hours worked. Two, this effect is mainly due to the extensive margin. Three, the effect is concentrated among less educated individuals and those already in poor health. Four, rises in frailty also reduces wages of workers without a college degree. Five, there is no evidence of significant effects of changes in earnings on health.

These findings suggest that health inequality may be an important source of lifetime earnings inequality. Moreover, the fact that the impact works primarily through the participation margin, suggests that the SSDI program may be an important factor driving this effect. Highly frail individuals have an increased likelihood of obtaining SSDI benefits and SSDI application and recipiency generates strong work disincentives. Individuals who apply for SSDI must be unemployed or have very low earnings, lower than the substantial gainful activity (SGA) threshold, for at least 5 months before benefit receipt can occur.[20] Once on DI, recipients whose earnings rise above the SGA threshold face the risk of losing benefits. As a result, most SSDI beneficiaries do not work.[21]

3 The Model

Given the findings in Section 2, we build a structural model that features individuals with risky and heterogeneous frailty profiles. We focus the model on the participation margin: individuals chose to participate in the labor market or exit and apply for SSDI. Given that we did not find any statistically significant effects of frailty on hours conditional on working, we do not model the intensive margin of labor supply. However, we do allow for poor health to impact individual's labor productivity and for this effect to be concentrated in individuals with less education. Since we failed to find a statistically significant feedback effect from earnings to health, we do not include one in the model.

3.1 Demographics

Time is discrete and one period is one year. The economy is populated by a continuum of individuals in J overlapping generations. The population grows at rate ν. Each period an age $j = 1$ cohort is born and lives up to the maximum age $j = J$. Individuals' health status

[19] This, of course, does not rule out the possibility that effects may still exist for specifically chosen subgroups.

[20] For example, in 2019 the SGA threshold was $1,220 per month.

[21] Maestas et al. (2013) document that in 2005 only 13 percent of SSDI applicants who started receiving benefits three years ago were employed.

is summarized by their frailty index, f, which evolves stochastically as we describe below. Frailty affects labor productivity, disutility from working, out-of-pocket medical expenditures and mortality risk. It also affects the chance of becoming a SSDI beneficiary. At each age j, the probability of surviving one more year depends on frailty, f, and education level, s, and is denoted by $p(j, f, s)$. Individuals are ex ante heterogeneous with respect to their education level, and they face a labor productivity process that is uncertain due to its dependence on both their frailty and direct labor productivity shocks.

Individuals derive utility from consumption, c, and (if working) suffer disutility from work which depends on frailty, f. Before retirement, each individual is either employed, non-employed or enrolled in Disability Insurance (DI). An employed individual works a fixed (exogenously given) fraction of time and earns wage $w \cdot \eta(j, f, s, \epsilon)$ which is the product of two terms. The first term is the wage per efficient unit of labor services, w. The second term is the efficiency unit of labor services per hour worked, $\eta(j, f, s, \epsilon)$. This depends on the worker's age j, frailty f, education s, and a stochastic component ϵ. The stochastic component, ϵ, consists of both a fixed effect and a persistent shock.[22] It evolves according to transition probability $\pi^e(\epsilon'|j, \epsilon, s)$, which depends on age j and education s. Employed workers may choose to quit and become non-employed. They can also become exogenously separated from their job with probability σ.

A non-employed individual can apply for DI or choose to go back to work immediately. If he applies for DI, he is awarded benefits with probability $\theta(f, n_a)$ in the next period. Here, n_a indicates the number of times he has applied for DI consecutively in the past. Individuals who are awarded DI benefits remain on DI until age $R < J$. After that, they transition to receiving social security retirement benefits. Those who choose to go back to work, have to pay a penalty $\chi(w\eta)$, which is a function of their current wage and can be understood of as the cost of job search.[23]

Those who are older than retirement age R receive social security retirement benefits but can choose to work or retire. Once an individual chooses to retire he remains retired until death. Both social security retirement and social security disability benefits are given by $SS(\bar{e})$, which is a function of the beneficiary's past earning history, $\bar{e}$.

Everyone has access to a risk-free asset a that pays return r. There are no other financial assets in the economy.

3.2 Frailty and medical expenditures

An individual's frailty is given by $f \equiv \psi(j, s, \epsilon_f)$. It depends on his age, j, education level, s, and a stochastic component, ϵ_f. The stochastic component consists of a fixed effect, a persistent shock, and a transitory shock. An individual whose stochastic component of frailty this period is given by ϵ_f will have value ϵ_f' next period with probability $\pi^f\left(\epsilon_f'|j, \epsilon_f, s\right)$.

[22] We do not include a transitory shock directly in the productivity process. However, the fact that individuals in the model face a positive probability of an exogenous job separation means that there is a transitory component to earnings risk.

[23] We do not explicitly model unemployment and job search. This modeling choice is motivated by the fact that the average duration of unemployment in the US is 15–20 weeks which is shorter than a period in our model. Therefore, we only include the monetary/income costs of short-term joblessness and abstract from the details of unemployment and job search.

Out-of-pocket medical expenditures are a deterministic function of age, education, frailty, and employment status. An individual of age j and education s who has frailty f incurs out-of-pocket medical expenditures $m^i(j, f, s)$ where $i = E, N, D, R$ depending on whether he is employed (E), non-employed (N), a DI beneficiary (D) or retired (R).[24]

3.3 Government

The government makes three distinct transfers to individuals that depend on the individuals' state:

- Social Security (SS): individuals aged R and older with earning history $\bar{e}$ receive social security benefit $SS(\bar{e})$ regardless of whether or not they are working.
- Disability insurance (DI): a DI beneficiary with earning history $\bar{e}$ receives DI benefit $SS(\bar{e})$.[25]
- Means-tested transfers: individuals with assets, a, and after-tax income net of medical expenditures and job search costs, y, receive transfer $Tr(a, y)$. The transfer is zero if $a + y \geq \underline{c}$. Otherwise, it is just enough to provide a minimum level of consumption of $\underline{c}$.

In addition, the government has exogenous expenditures G. To finance these expenditures and transfers, the government levies a nonlinear tax on labor income, $T(w\eta)$, and a proportional tax on capital income, τ_K (paid by the firm). Due to the absence of annuity markets, individuals may die with positive assets. We assume that these accidental bequests are taxed at a rate of 100 percent by the government.

3.4 Individual decision problems

To economize on notation we denote a subset of the state space as $x \equiv (j, a, f, s, \epsilon, \bar{e})$.[26] Let $V^E(x, i_s)$ be the value function of an employed individual, $V^N(x, n_a)$ be the value function of a non-employed individual, $V^D(x, n_d)$ be the value function of a DI beneficiary, and $V^R(x)$ be the value function of a retiree. The variable i_s is an indicator that an employed worker is returning from an exogenous separation or non-employment spell. Variable n_a tracks the number of periods an individual has been in non-employment consecutively in the past. Recall that workers can always go back to employment immediately. If they stay in non-employment, it is because they are applying for DI benefits. Therefore, n_a is also equal to the number of times an individual has applied for DI consecutively in the past. Variable n_d represents the number of periods an individual has been on DI. This variable is used to determine his eligibility for Medicare benefits. Individuals discount the future at rate β. We

[24] All workers who are older than age R are Medicare beneficiaries and face the same process for out-of-pocket medical expenditures as retirees.

[25] The U.S. Social Security administration uses the same benefit formula to calculate both retirement and disability benefits.

[26] To avoid clutter, we use x as the argument of functions with the understanding that not all functions depend on all the elements of vector x.

now describe the problems facing each type of individual.

The employed worker's problem: employed workers face the risk of exogenously separating from their employer with probability σ at the beginning of the next period. If separated, they can choose to go back to work immediately. If they survive the separation shock, they can choose to quit the job voluntarily. When $j < R-1$, their utility-maximization problem can be specified as follows,

$$\begin{aligned} V^{E}\left(x, i_s\right) = & \max_{c, a' \geq 0} u\left(c, v\left(f\right)\right) \\ & + \sigma \beta p\left(j, f, s\right) E\left[\max\left\{V^{E}\left(x', 1\right), V^{N}\left(x', 0\right)\right\}\right] \\ & + \left(1-\sigma\right) \beta p\left(j, f, s\right) E\left[\max\left\{V^{E}\left(x', 0\right), V^{N}\left(x', 0\right)\right\}\right] \end{aligned} \tag{6}$$

subject to

$$\frac{a'}{1+r} + c + m^{E}\left(j, f, s\right) = a + w\eta\left(x\right) - T\left(w\eta\left(x\right)\right) - \chi\left(w\eta\left(x\right)\right) i_s + Tr\left(x, i_s\right), \tag{7}$$

and $\bar{e}' = [(j-1)\,\bar{e} + w\eta\left(x\right)]/j$.

When workers return from a separation or non-employment ($i_s = 1$), they have to pay a penalty $\chi\left(w\eta\left(x\right)\right)$, which is a function of their hourly wages. Employment decisions of these workers at the beginning of the period are denoted by $I_E(x, i_s)$.

After reaching age R, employed workers can only choose between working and retirement (not working). However, they are eligible to claim social security retirement benefits regardless of whether they work or not.[27] They are also eligible for Medicare, which affects their out-of-pocket medical expenditures. Therefore, employed workers of age $j \geq R-1$ face the following optimization problem,

$$\begin{aligned} V^{E}\left(x, i_s\right) = & \max_{c, a' \geq 0} u\left(c, v\left(f\right)\right) \\ & + \sigma \beta p\left(j, f, s\right) E\left[\max\left\{V^{E}\left(x', 1\right), V^{R}\left(x'\right)\right\}\right] \\ & + \left(1-\sigma\right) \beta p\left(j, f, s\right) E\left[\max\left\{V^{E}\left(x', 0\right), V^{R}\left(x'\right)\right\}\right] \end{aligned} \tag{8}$$

subject to

$$\begin{aligned} \frac{a'}{1+r} + c + m^{R}\left(j, f, s\right) &= a + w\eta\left(x\right) + SS(\bar{e}) - T\left(w\eta\left(x\right)\right) - \chi\left(w\eta\left(x\right)\right) i_s + Tr\left(x, i_s\right), \quad (9) \\ \bar{e}' &= \bar{e}. \end{aligned}$$

The non-employed's problem: non-employed individuals apply for DI. They qualify for benefits with probability $\theta\left(f, n_a\right)$. If awarded, they start receiving them in the following

[27] We make this assumption for simplicity. Full retirement age is by far the most common claiming age in the data and, aside from some tax implications, which Jones and Li (2018) find to have a relatively small effect on the labor supply of older workers, there is no cost of working past full retirement age while also claiming benefits.

period and will remain on DI until they reach retirement age R.[28] At that time, they move to social security. If not awarded, they can go back to work immediately or remain non-employed and apply again. When $j < R-1$, the non-employed individual's problem can be specified as follows,

$$\begin{aligned} V^N(x, n_a) = & \max_{c,a' \geq 0} u(c) \\ & + \theta(f, n_a)\, \beta p(j, f, s)\, E\left[V^D(x', 0)\right] \\ & + (1 - \theta(f, n_a))\, \beta p(j, f, s)\, E\left[\max\left\{V^E(x', 1), V^N(x', n_a + 1)\right\}\right] \end{aligned} \tag{10}$$

subject to

$$\frac{a'}{1+r} + c + m^N(j, f, s) = a + Tr(x, n_a). \tag{11}$$

Employment decisions of these workers are denoted by $I_N(x, n_a)$.

When $j = R - 1$, non-employed individuals cannot apply for DI anymore as they will reach the retirement age in the next period. The problem facing them becomes,

$$V^N(x, n_a) = \max_{c,a' \geq 0} u(c) + \beta p(j, f, s)\, E\left[\max\left\{V^E(x', 1), V^R(x')\right\}\right] \tag{12}$$

subject to (11).

The DI beneficiary's problem: DI recipients only make consumption and saving decisions. It is important to note that DI recipients can also get access to Medicare benefits after enrolled in DI for two years. In the model, this eligibility is determined by the state variable n_d, which represents the number of periods the individual has been on DI. When $j < R - 1$, DI recipients face the following problem,

$$V^D(x, n_d) = \max_{c,a' \geq 0} u(c) + \beta p(j, f, s)\, E\left[V^D(x', n_d + 1)\right] \tag{13}$$

subject to

$$\frac{a'}{1+r} + c + m^D(j, f, s, n_d) = a + SS(\bar{e}) + Tr(x, n_d). \tag{14}$$

When DI beneficiaries reach retirement age R, they automatically move from disability insurance to social security. Therefore, for $j = R - 1$,

$$V^D(x, n_d) = \max_{c,a' \geq 0} u(c) + \beta p(j, f, s)\, E\left[V^R(x')\right] \tag{15}$$

subject to (14).

The retiree's problem: retirees remain retired until they die. They receive social security benefits and only make consumption and saving decisions. Their problem is given by

[28] We do not model exits from DI due to reasons other than transition to old-age social security or death because they are extremely rare. According to the Social Security Administration, in 2018, the fraction who exited due to the next two most common reasons were 0.6 percent (who exited because they earned more than the maximum allowed level) and 0.5 percent (who exited because they were deemed medically able to work during a medical review).

$$V^R(x) = \max_{c,a' \geq 0} u(c) + \beta p(j,f,s) E\left[V^R(x')\right] \tag{16}$$

subject to

$$\frac{a'}{1+r} + c + m^R(j,f,s) = a + SS(\bar{e}) + Tr(x). \tag{17}$$

3.5 Technology

There is a representative firm that produces a single good using a Cobb-Douglas production function such that $Y = AK^{\alpha}N^{1-\alpha}$ where α is the output share of capital, K and L are the aggregate capital and aggregate labor input, and A is the total factor productivity. Capital depreciates at a constant rate $\delta \in (0,1)$. The firm pays a proportional tax on capital income τ_k. We assume a small open economy such that the after-tax return on assets, r, is exogenous. Therefore, in equilibrium, capital per worker is given by

$$r = (1-\tau_k)\left(\alpha A(K/N)^{\alpha-1} - \delta\right),$$

which determines the wage per efficient unit of labor services,

$$w = (1-\alpha)A(K/N)^{\alpha}. \tag{18}$$

We assume that the economy is in a stationary competitive equilibrium. The full definition of the equilibrium is provided in Section 3 of the Online Appendix.

4 Calibration

Our calibration strategy consists of two stages. In the first stage, we set the values of some parameters that can be determined based on independent estimates from the data or the existing literature. In the second stage, we calibrate the rest of the parameters by minimizing the distance between data targets and their model counterparts.

Our goal is to quantify the impact of health inequality on lifetime earnings inequality. To do so we must first pin down the magnitudes of the various channels through which frailty impacts earnings and employment in the model. Recall that the five channels through which frailty operates are via its impact on: 1) mortality rates, 2) out of pocket medical expenditures, 3) labor productivity, 4) probability of successful DI application, and 5) disutility from working. The effect of frailty on the first three channels can be estimated directly from the data without using the model. As we describe in more detail below, we estimate these effects in the first stage of the calibration.

The effects of frailty on the probability of successful DI application and the disutility from working cannot be discerned directly from the data.[29] These effects are, instead, determined in the second stage of the calibration by minimizing the distance between model and data moments. Specifically, the parameters that determine the frailty effects together with the

[29]We cannot directly estimate the probabilities of successful DI application because none of the datasets we use provides information on whether or not a respondent has applied for DI. We only see whether or not respondents are currently receiving DI benefits.

other parameters governing DI eligibility and disutility from work are chosen by targeting labor force participation rates and SSDI recipiency rates by age and frailty percentile groups. The moments are concentrated in the unhealthy tail of the frailty distribution since this is where the effects of frailty on labor supply and SSDI recipiency are most pronounced.

The targeted rates are shown in the left and right panels of Figure 3.[30] Notice that this set of target moments includes labor force participation rates by frailty of both younger workers and workers over the age of 65. This is intentional. The chance of successful DI application does not directly impact the labor supply choices of individuals after the age of 65. Thus, by targeting the dispersion in labor force participation rates by frailty for this age group we are able to separately identify frailty's impact on the disutility from work from its impact on successful DI application.

4.1 Demographics and initial distributions

We assume age $j = 1$ corresponds to a 25 year-old and $J = 70$ corresponds to a 94 year-old. Workers receive old-age social security and Medicare benefits at age $R = 41$ (66 year-olds). This is also the age at which they no longer choose between working and applying for SSDI but instead choose between working and being retired.

Conditional survival probabilities at each age are estimated using HRS data and a probit regression. Mortality depends on a quadratic in frailty, a quadratic in age, education, and gender. The results of this estimation are presented in Section 4 of the Online Appendix.[31] We adjust the value of the estimated constant so that population mortality is consistent with the year 2000 period life-table in Bell and Miller (2005). The population growth rate is set to $\nu = 0.02$ so that the ratio of old (over 65) to young (65 and younger) is equal to 0.2 (this is consistent with the year 2000 U.S. Census).

The population is divided into three education groups: high school dropouts, high school graduates, and college graduates. The initial distribution of agents across the three groups is 12 percent high school dropouts, 52 percent high school graduates, and 36 percent college graduates based on the education distribution of 25–26 year-olds in our PSID sample.

Even though the fraction of individuals non-employed and on DI is low at ages 24–26, it varies substantially across frailty and education. For this reason, we set the initial distributions of individuals across employment states (employed, non-employed, and DI beneficiary) by education and frailty percentile group to be consistent with their counterparts in the data. Section 4 of the Online Appendix provides the numbers.

[30] The left panel of the figure is constructed using our PSID sample. Individuals are considered employed if they work at least 260 hours per year and earn at least \$3 per hour. The right panel is constructed using MEPS data. MEPS does not contain information on DI beneficiary status. However, it does include data on whether an individual receives Medicare benefits. DI beneficiaries are the only group younger than 65 years of age who qualify for Medicare (after being on DI for two years). We compute the fraction in each frailty and age group who receive Medicare benefits in MEPS. We then adjust the fractions such that we replicate DI recipiency rates by age in the population provided by the Social Security Administration.

[31] Agents do not have a gender in the model. Instead, we compute the mortality rate of an agent by giving him the average gender in the population in year 2000.

Table 4: Estimated effects of one additional frailty deficit on log productivity (wage)

	no bias correction	with bias correction
$\text{frailty}_t \times$ HSD	-0.042**	-0.044**
	(0.017)	(0.017)
$\text{frailty}_t \times$ HS	-0.025***	-0.027***
	(0.009)	(0.009)
$\text{frailty}_t \times$ CL	-0.002	-0.001
	(0.004)	(0.004)

Note: Standard errors are in parenthesis. $*p < 0.1$; $**p < 0.05$; $***p < 0.01$.

4.2 Preferences

Individuals have utility over consumption, c, and suffer disutility from working that depends on their frailty, f. The period utility is given by,

$$u(c, f) = \frac{\left(c^{\mu}\left(1 - v(f) \times i_p\right)^{1-\mu}\right)^{1-\gamma}}{1-\gamma},$$

with

$$v(f) = \phi_0 \left(1 + \phi_1 f^{\phi_2}\right),$$

where $i_p = 1$ if the individual is working and 0 otherwise. For the benchmark calibration, we set $\gamma = 2$ and $\mu = 0.5$, which implies a coefficient of relative risk aversion of $1 - (1-\gamma)\mu = 1.5$. This is in the middle of the range of values used in the macro literature.[32]

We assume $\phi_0 \geq 0$, $\phi_1 \geq 0$, and $\phi_2 \geq 0$ so that higher levels of frailty increase the disutility of working. If $\phi_2 > 1$ then $v(f)$ is convex in frailty and the marginal effect of increasing frailty is higher for more frail individuals. The opposite is true if $\phi_2 < 1$. As we explained above, the parameters ϕ_0, ϕ_1, and ϕ_2 are determined in the second-stage minimization. They are pinned down by the variation by age and frailty in the labor force participation rates of workers aged 25 to 84.

4.3 Labor productivity and job separation

We estimate the labor productivity process, $\eta(j, f, s, \epsilon)$, separately for each education group using PSID data. For each group, labor productivity is the sum of a deterministic component and a stochastic component. The deterministic component consists of a cubic in age and a linear frailty effect. The stochastic component contains both a fixed effect and an AR(1) shock.

One concern when estimating the labor productivity process is selection bias. We do not observe hourly wages (our proxy for labor productivity) of those who do not work. If individuals whose frailty more negatively impacts their labor productivity are less likely to work, not controlling for selection will lead us to underestimate the impact of frailty on productivity.

[32] See Attanasio (1999) and Blundell and MaCurdy (1999) for surveys.

To correct for potential selection bias, we estimate the labor productivity process in three steps. First, we use the system GMM dynamic panel estimator outlined in Section 2.2 and a selection correction procedure to estimate the effect of frailty on productivity. Second, removing the frailty effects from our productivity observations, we estimate the age effects via OLS. Third, using variance-covariance moments constructed with the final frailty residuals, we estimate the stochastic component via GMM.

To conduct the selection correction in the first step, we follow Al-Sadoon et al. (2019) who show that in system GMM, selection bias is mainly due to correlation of the fixed effects in the selection and outcome processes. To correct for this selection bias they propose first estimating a selection equation that includes fixed effects and an exclusion restriction.[33] They then suggest including the estimated fixed effects as regressors in the outcome equation.

We use a fixed effect linear probability model of employment as our selection equation. We include the same set of regressors in our selection equation as in the outcome equation with the addition of exclusion restrictions. Following Low and Pistaferri (2015), we use "potential" government transfers interacted with frailty as our exclusion restrictions. "Potential" government transfers are defined as the sum of food stamps, AFDC/TANF payments, unemployment insurance benefits, and EITC payments that individuals would receive if they applied. These transfers vary across states. Moreover, they depend on marital status and number of kids (which also vary across individuals in our sample). Therefore, instead of using the amount of "potential" transfers as exclusion restrictions we use the interaction of state of residence, number of kids and martial status (a total of 482 combinations). These "potential" transfers do not directly impact individuals wages or labor productivity but do create different work incentives for people with different frailty levels.

Once we have estimates of the first-stage fixed effects, we estimate the effect of frailty on log wages using the dynamic panel system GMM estimator. For each individual at each age, log wages is assumed to be a function of two lags of log wages, lagged frailty interacted with education, a fixed effect, and the individual's estimated fixed effect from the first stage regression. Table 4 reports the estimated effects of accumulating one additional deficit on log wages for each education group with and without controlling for selection.[34][35] Notice that, consistent with our concerns, the effect of frailty on wages is slightly smaller when not controlling for selection bias. Also notice that, consistent with our findings in Section 2.2, the negative effects of frailty on wages are decreasing with education and only significant for workers without a college degree. One additional deficit reduces wages by 4.4 percent for high school dropouts, 2.7 percent for high school graduates and less than 0.1 percent for college graduates.[36]

[33] An exclusion restriction is a variable that impacts the selection process but not the outcome process. In our case, a valid exclusion restriction should impact employment but not labor productivity or wages.

[34] The first-stage and full set of second-stage estimation results can be found in Section 4 of the Online Appendix.

[35] Note two important differences between the estimation in this section and in Section 2.2. First, the estimation in this section is done only on a sample of men. Second, we treat frailty as exogenous in this regression given our earlier finding on the absence of reverse causality. The estimated effects of frailty on productivity are robust to making frailty endogenous.

[36] Using a different measure of health to ours, Low and Pistaferri (2015) also estimate the effect of poor health on labor productivity. In a rough comparison, we find that our effects are similar in magnitude, albeit slightly larger, as compared to theirs. Details are provided in Section 4 of the Online Appendix.

Given these estimation results, we next remove the effects of frailty from our log wage observations. Then, to obtain the deterministic age effects, separately for each education group, we regress the adjusted log wages on a cubic polynomial in age and year dummies. Finally, we use the residuals from this regression to construct a set of variance-covariance moments. As in Guvenen (2009), we use these moments to estimate the stochastic component via GMM. For these last two steps, we group together high school dropouts and high school graduates in order to take advantage of larger sample sizes. The implication of this is that, while the effect of frailty on productivity is education group specific, the age effects and stochastic component are the same for both high school dropouts and high school graduates. The estimation results are presented in Section 4 of the Online Appendix.

Finally, the job separation rate is set to $\sigma = 15\%$, which is the average (annual) rate of layoffs and discharges between 2005 and 2007 according to the Jobs Opening and Labor Turnover Survey (JOLTS) by the Bureau of Labor Statistics (BLS).

4.4 Frailty and medical expenditures

Our specification and estimation of the frailty process, $\psi(j, s, \varepsilon_f)$, using PSID data follows closely that in Hosseini et al. (2019). We assume that there is a positive mass of individuals with zero frailty at age 25. Each period, these individuals move to a positive frailty value with a probability that depends on their education and a quadratic in age. Once positive, an individual's frailty never goes back to zero.[37] We use a probit regression to estimate the conditional probabilities of positive frailty by age and education.[38]

For individuals with positive frailty, log frailty is given by the sum of a quartic age polynomial and a stochastic component. The stochastic component consists of an AR(1) shock, a transitory shock, and a fixed effect. The AR(1) shock captures persistent health events such as developing diabetes, while the transitory shock captures acute ones such as a temporary inability to walk due to a broken leg. We find that there are large differences in frailty dynamics by education. For this reason, we estimate the log frailty process separately for each education group.

Frailty and mortality are highly correlated. Thus, when estimating the nonzero frailty process, it is important to control for selection bias due to mortality. To this end, we estimate the frailty processes using an auxiliary simulation model and the method of simulated moments (MSM). The auxiliary simulation model simulates the frailty dynamics described above together with the mortality rates by age and education given by the specification in Section 4.1. For each education group, the coefficients of the age polynomial are determined by targeting the age profile of log frailty for 25 to 95 year-old PSID respondents. The variance and persistence of the AR(1) shock, variance of the transitory shock, and variance of the fixed effect are determined by targeting variance-covariance moments by age of the log frailty residuals.[39]

Figure 2 shows the estimation results for high school graduates (the largest education group in our sample).[40] The left panel shows the fraction of high school graduates with zero

[37] Less than 1 percent of individuals in our PSID sample with positive frailty have zero frailty next period.

[38] The estimated parameters are reported in Section 4 of the Online Appendix.

[39] The estimated parameters are reported in Section 4 of the Online Appendix.

[40] The estimation results for the other two education groups are show in Section 4 of the Online Appendix.

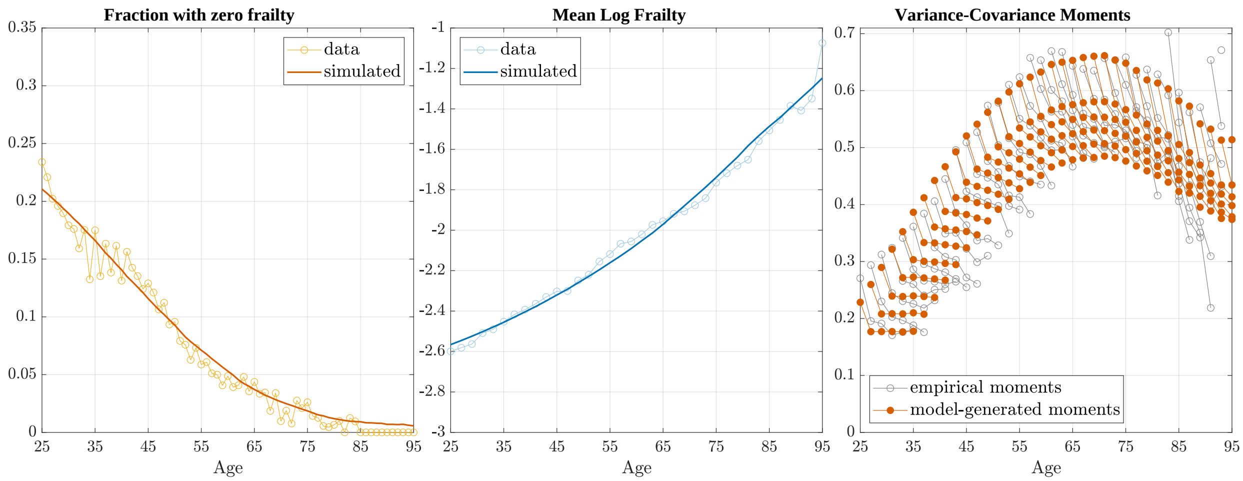

Figure 2: Estimation targets: auxiliary simulation model vs PSID data for high school graduates. Left panel is the fraction with zero frailty by age, middle panel is mean log frailty by age for those with nonzero frailty, and right panel is the age-profile of the variance and covariances of log frailty residuals (the stochastic component of log frailty).

frailty by age in the data and in the simulation of the model. The middle figure shows the age profile of mean log frailty targeted in the data and the model counterpart. The right panel shows the age profile of the variance-covariance moments in the model and the data. Notice that our estimated frailty process is able to generate autocovariance patterns that are very similar to those in the data.

We estimate out-of-pocket medical expenditures separately by education and labor market status: employed, non-employed, and on Medicare (which includes both retirees and those who are on DI). To capture the nonlinear effect of frailty on medical spending, we assume that log out-of-pocket medical expenditures are determined by a cubic in age and a cubic in frailty. We estimate the coefficients of these functions using data from MEPS. Note that although we do not include any randomness directly in this formulation, the out-of-pocket medical expenditure is random through its dependence on frailty. The results of these estimations are presented in Section 4 of the Online Appendix.

4.5 DI application

The Social Security Disability Insurance program (SSDI) application process is complex and lengthy. Moreover, even though the probability of successfully obtaining benefits is generally higher for individuals in worse health, the outcome is uncertain.[41] The process starts with a 5 month waiting period during which applicants are not allowed to be gainfully employed. After this initial period, applicants' cases are reviewed by the Disability Determination Service review board. The most definite cases are approved for benefits at this point. For instance, individuals with one of 100 specifically listed and verifiable medical conditions are usually given benefits at this stage. Less definite cases are usually denied. However,

[41] See French and Song (2014) and the references therein for a detailed overview of the program.

Table 5: Parameters chosen outside the model

Parameter	Description	Values/source
Demographics		
J	maximum age	70 (94 y/o)
R	retirement/SS eligibility age	41 (66 y/o)
ν	population growth rate	0.02
Preferences		
γ	curvature of utility function	2
μ	weight on consumption	0.5
Job Separation		
σ	annual layoffs/separations in JOLTS	0.15
Technology		
α	capital share	0.33
δ	depreciation rate	0.07
r	return on assets	0.04
Government policies		
τ_{SS}, τ_{med}	social security and Medicare tax rates	0.124, 0.029
τ_K	capital tax	0.3
τ	tax progressivity	0.036
$\underline{c}$	minimum consumption (% of ave. earnings)	11
G	government purchases (% of GDP)	17.5

after a 60 day waiting period, denials can be appealed. Such appeals are assessed by a judge after a period of roughly one year. Judges have considerable latitude in assessing appeals. Applicants whose appeals are denied, can continue to appeal for multiple rounds with approximately a one year turnaround time between appeal and decision each round. Alternatively, denied applicants can end the appeals process and start over applying for benefits by submitting a new application.

French and Song (2014) document that by one year after initial application, about 50 percent of applicants will usually have been awarded benefits. After this point, the probability of obtaining benefits continually declines in the number of years since initial application (see the middle panel of Figure 3). After 10 years, only 70 percent of applicants who continually appeal or reapply are successful. Thus, individuals can spend years trying to successfully get on DI.

Motivated by the description above, we assume that the probability of successful DI application depends on an individual's frailty, f, and the number of periods the individual has subsequently applied for DI, n_a. Specifically, we set

$$\theta\left(f, n_a\right) = \min\left\{1, \kappa_0 f^{\kappa_1} n_a^{\kappa_2}\right\}.$$

The parameters κ_0, κ_1, and κ_2 are determined in the second stage of the calibration. Both κ_0 and κ_1 are identified by targeting DI recipiency rates by age and frailty in the minimization. Parameter κ_2 is identified off the rate at which the probability of obtaining benefits declines in the years since initial application. In fact, we include n_a in the probability function to

Table 6: Parameters calibrated using the model

Parameters	Description		Values
Preferences			
β	discount factor		0.982
Policy			
λ	HSV tax parameter(level)		0.908
Disutility of work			
ϕ_0	level		0.636
ϕ_1	frailty effect		1.2
ϕ_2	frailty effect		3.0
Prob. of successful DI application			
κ_0	level		50
κ_1	frailty effect		5.0
κ_2	number of attempts effect		-0.1
Targeted Moments		Data	Model
Wealth-earning ratio		3.2	3.2
Federal income tax (% of GDP)		8	8
LFP rates by frailty percentile group and age		Figure 3	
DI recipiency rates by frailty percentile group and age		Figure 3	
Relative DI success rates		Figure 3	

insure that the model will match this rate. Matching this rate is important as it impacts the decision to continue to apply for DI or reenter the labor force. As Table 6 shows, however, the minimization procedure places little weight on n_A. This is because the probability of successfully obtain DI is increasing and convex in frailty. As a result, the applicant pool naturally becomes more and more selected towards healthier, and hence less likely to be successful, applicants as the number of periods since initial application increases.

4.6 Policy parameters

For old-age social security and DI benefits, we use the social security benefit formula for primary insurance amount (PIA):

$$SS\left(\bar{e}\right)=\begin{cases}0.9\bar{e} & \bar{e}\leq 0.2\bar{e}_a,\\ 0.18\bar{e}_a+0.33\left(\bar{e}-0.2\bar{e}_a\right) & 0.2\bar{e}_a<\bar{e}\leq 1.25\bar{e}_a\ ,\\ 0.5265\bar{e}_a+0.12\left(\bar{e}-1.25\bar{e}_a\right) & 1.25\bar{e}_a<\bar{e},\end{cases}$$

where $\bar{e}_a$ is the average earnings in the economy.

The tax function $T(\cdot)$ has three components. One is a nonlinear component mimicking the U.S.income tax/transfer system. One is a social security payroll tax component consisting of a proportional tax that is subject to a maximum taxable earnings cap. And, one is a proportional Medicare payroll tax. We model the nonlinear component of the tax function in the fashion of Benabou (2002) and Heathcote et al. (2017). That is, the tax function $T(\cdot)$

is given as follows,

$$T(e) = e - \lambda e^{1-\tau} + \tau_{ss} \min\{e, 2.47\bar{e}_a\} + \tau_{med} e.$$

Here, τ controls the progressivity of the tax function and is set to 0.036 based on the estimate by Guner et al. (2014). We choose the value of λ in the second-stage of the calibration to match the total federal income tax receipts as a share of GDP in the U.S. data.

The social security payroll tax rate is set to $\tau_{ss} = 0.124$ and the Medicare tax rate is set to $\tau_{med} = 0.029$. There is also a capital tax, $\tau_K = 0.3$, which is paid by the firm and set based on Gomme and Rupert (2007). The minimum consumption level, c_{min}, is set to 11 percent of average earnings.[42] Finally, we set exogenous government purchases, G, to 17.5 percent of GDP. We hold this ratio fixed in all counterfactul experiments we run.

4.7 Technology

We assume a small open economy and set $r = 0.04$. The capital share α is set to 0.36. We normalize aggregate TFP, given by A, to 1, and choose β in the second stage such that the model generates a wealth-to-earnings ratio of 3.2.[43] The depreciation rate is set to $\delta = 0.07$. This is based on calculations in Gomme and Rupert (2007).[44]

For employed workers who just came back from non-employment, we assume that they suffer a wage penalty which mimics the forgone earnings during job search within the period. According to the U.S. Bureau of Labor Statistics, the average duration of unemployment in the U.S. was approximately 15–20 weeks between 2000 and 2007. Therefore, we set the wage penalty to be a third of one year's earnings.

5 Assessment

Tables 5 and 6 summarize the baseline model parameterization. Figure 3 provides a comparison of the labor force participation rates, DI recipiency rates, and relative DI success rates targeted in the data with the model counterparts. All targeted moments are reasonably matched. Notice that although the model slightly understates the labor force participation rates of workers ages 75–84, it matches well the level and dispersion in labor force participation rates of workers aged 65–74. This is important as it is these rates relative to the rates of those under 65 that determines the disutility from work parameters.

To assess the model's performance with regards to non-targeted moments, we look at labor force participation and DI recipiency rates across different education groups. Table 7 shows the overall participation and DI recipiency rates by education group. The top panel shows that the model does reasonably well in matching overall participation rates by

[42] See Kopecky and Koreshkova (2014) for details and a review of the literature.

[43] In choosing this moment we follow Hong and Ríos-Rull (2012) and only target the wealth-to-earnings ratio of the bottom 95 percent.

[44] Gomme and Rupert (2007) calculate depreciation rates for four different sectors (market structures, equipment and software, housing, and consumer durables). We use a weighted average of their calculations. To do this we weight depreciation in each sector by the shares of capital in that sector as reported in their paper.

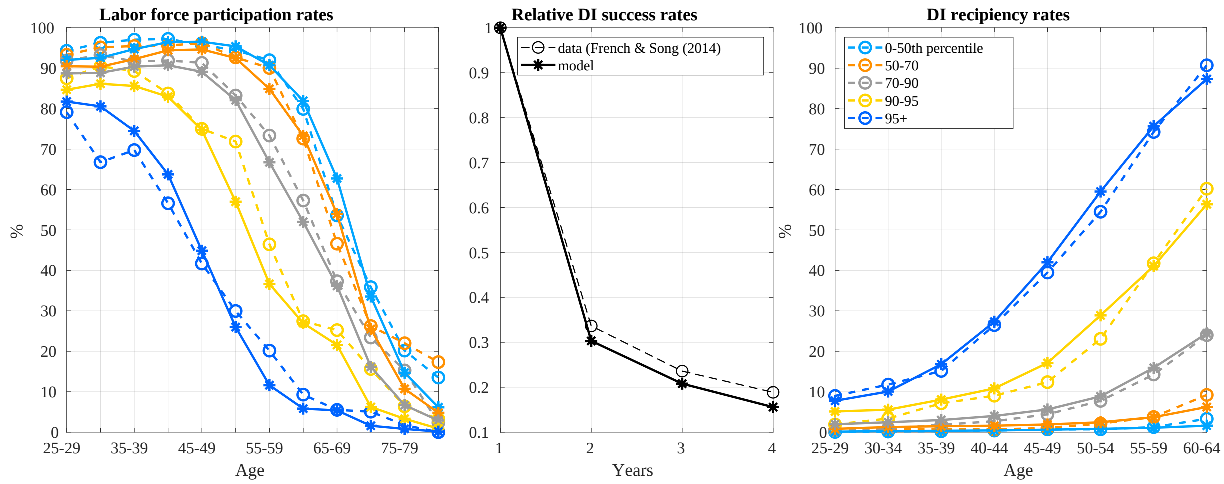

Figure 3: Calibration targets: model versus data. Solid lines are the model. Dashed lines are the data. The left panel shows the labor force participation rates (source for data is PSID), the right panel shows the DI recipiency rates (source for data is MEPS), and the middle panel shows the relative rates of success in SSDI application by years since initially applied (data moments are based on findings reported by French and Song (2014)). The rate for the first year is normalized to 1.

education. The bottom panel shows that the model slightly understates the SSDI recipiency rate for college graduates and slightly overstates it otherwise.

To inspect the fit by education further, we report the labor force participation and DI recipiency rates by age and frailty percentile groups separately for each education group. The labor force participation rates are presented in Figure 4 and the DI recipiency rates are presented in Figure 5. From these figures we see that the model performs reasonably well in capturing the pattern of labor force participation for each education group. In particular, it captures well the reduced dispersion in both labor force participation rates and DI recipiency rates by frailty as education increases. For those in the top 5 percent of the frailty distribution, the model slightly understates the labor force participation rates of the high school dropouts and slightly overstates the participation rates for college graduates. Likewise, the model overstates the DI recipiency rates of high school dropouts with frailty in the top 5 percent and understates the recipiency rates for college graduates in this group. Recall that in our dynamic panel estimation we did not find any effect of frailty on labor productivity for college graduates. However, it is possible that, even though there is no effect on average, there is a positive effect for those at the very top of the frailty distribution.[45] If this is in fact the case, adding the effect would improve the fit of the model by education for this group.

Overall, the model does an excellent job matching the dispersion in labor force participation and DI recipiency rates by age, frailty, and education. Given that this is the case, we now move on to the quantitative analysis of the impact of health inequality on lifetime

[45]We did explore this possibility by adding a quadratic frailty effect in the dynamic panel estimation described in Section 4.3 with mixed results. Ultimately, we were unable to find a specification that passed all model validity tests and indicated a statistically significant nonlinear frailty effect.

Table 7: Labor force participation rates (top) and DI recipiency rates (bottom) by education groups

	Labor force participation rates (%)		
	High School Dropouts	High School Graduates	College Graduates
Data	78	87	93
Model	77	86	94
	DI recipiency rates (%)		
	High School Dropouts	High School Graduates	College Graduates
Data	9.5	5.0	1.4
Model	10.3	5.8	1.0

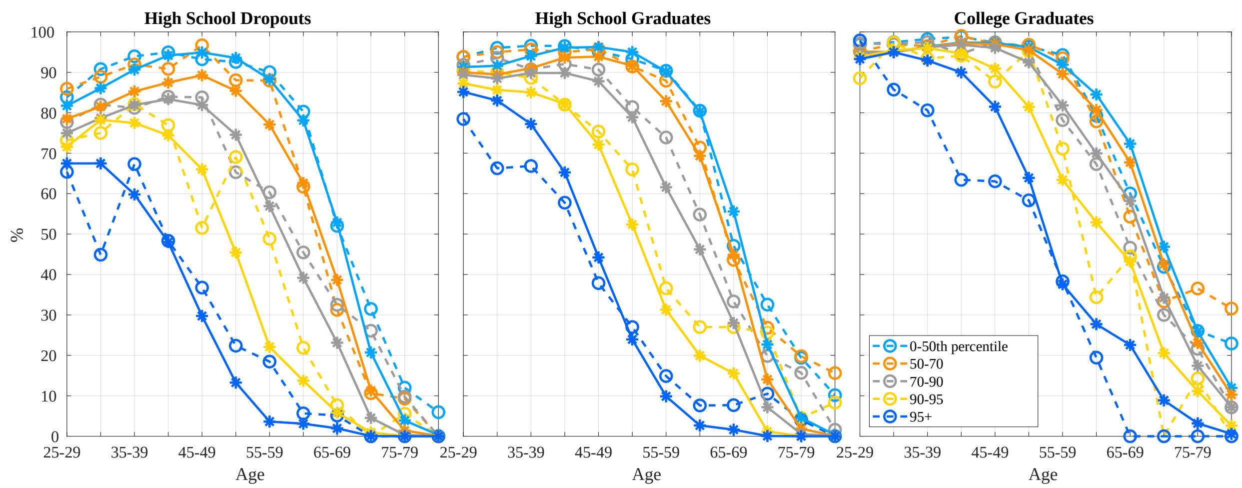

Figure 4: Calibration assessment: model versus data. Solid lines are the model. Dashed lines are the data. The panels show the labor force participation rates by frailty percentile groups and age for each education group: high school dropouts, high school graduates, and college graduates. Source for the data is PSID.

earnings inequality.

6 Quantitative Exercise

In this section we use the calibrated model to assess the impact of health inequality on lifetime earnings inequality.[46] To do this we conduct the following experiment. We consider a counterfactual economy in which everyone has the average frailty profile. Giving all individuals in the economy the average frailty profile removes all cross-sectional variation in frailty conditional on age. In particular, it removes heterogeneity in frailty due to education and individuals' fixed frailty types. It also removes the heterogeneity in frailty due to the

[46]We focus on lifetime earnings given that much of the impact of health on earnings comes via permanent reductions in labor force participation. Another paper which studies the lifetime earnings distributions in the U.S. is Guvenen et al. (2017).

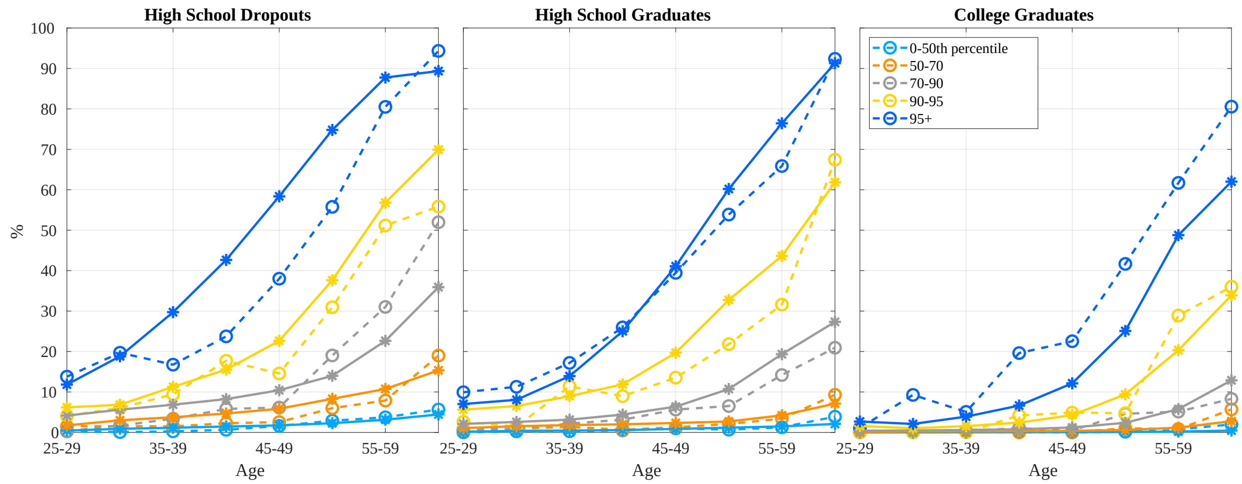

Figure 5: Calibration assessment: model versus data. Solid lines are the model. Dashed lines are the data. The panels show the DI recipiency rates by frailty percentile groups and age for each education group: high school dropouts, high school graduates, and college graduates. Source for the data is MEPS.

persistent and transitory frailty shocks. We refer to the counterfactual economy as the No-Frailty-Heterogeneity economy (or NFH for short). We compare the inequality in lifetime earnings at different ages between the NFH economy and the benchmark. For each individual, our measure of lifetime earnings at each age is simply that individual's accumulated earnings to date.

Removing health inequality significantly reduces inequality in lifetime earnings at older ages. Figure 6a shows the age-profile of the variance of log lifetime earnings in the benchmark economy and the NFH economy. The variation in lifetime earnings is almost the same in the two economies at younger ages. However, between ages 35 and 65, the variance of log lifetime earnings increases more rapidly with age in the benchmark economy. As a result, there is less variation in lifetime earnings in the NFH economy starting around age 40. As reported in Table 8, at age 45 the variance of log lifetime earnings is 12.9 percent lower in the NFH economy relative to the benchmark. The relative difference peaks at age 65 when the variance of log lifetime earnings is 28.9 percent lower.

Removing health inequality, not only reduces the variance of log lifetime earnings but it also leads to a smaller fraction of individuals with zero lifetime earnings at each age. As Figure 6b shows, the fraction of these individuals in the benchmark and NFH economy declines rapidly between ages 25 and 30 after which it remains small. Notice that, while small in both, the fraction of individuals with zero lifetime earnings is considerably lower in the NFH economy.[47]

Health inequality in the model is due to both initial heterogeneity in frailty, captured by fixed effects and education, and the effect of idiosyncratic frailty shocks. To understand the

[47] The lower fraction of 25 year-olds with zero lifetime earnings in the NFH economy is due to the fact that more individuals are initially employed in this economy. Recall that the initial distribution of individuals across employment states is conditional on frailty percentile groups. When everyone has the average frailty profile, the initial distributions are those of the 50–70th percentiles.

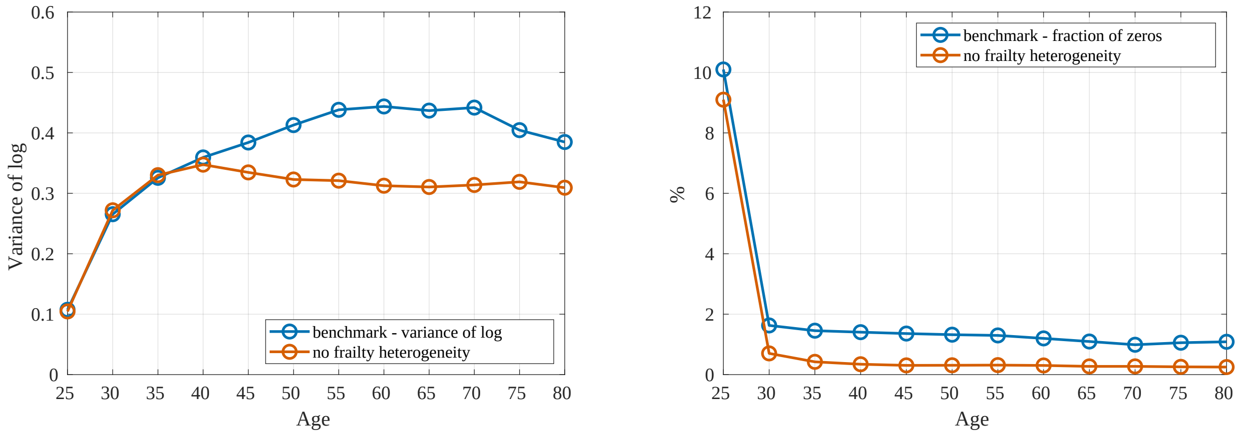

(a) Variance of log lifetime earnings.

(b) Fraction with zero lifetime earnings.

Figure 6: The variance of log lifetime earnings (left) and the fraction of individuals with zero lifetime earnings (right) in the benchmark economy (blue) and the No-Frailty-Heterogeneity economy (red).

relative importance of each of these components for lifetime earnings inequality, we conduct two related counterfactual experiments and report the results in Table 8. The middle panel of the table shows the variance of log lifetime earnings for a counterfactual economy that is identical to the benchmark, except that there are no frailty shocks. All inequality in frailty in this economy is due to initial fixed heterogeneity. The bottom panel shows the results of the same calculation for an economy that does not feature any initial heterogeneity in frailty due to fixed effects or education. All the inequality in frailty in this counterfactual is driven by frailty shocks. As the table shows, frailty shocks have a larger impact on lifetime earnings inequality than initial frailty heterogeneity. While frailty shocks are more important at all four ages reported in the table, their relative importance peaks at age 55 when their effect on lifetime earnings inequality is more than double. The variance of log lifetime earnings at age 55 is 10 percent lower in the economy without initial frailty heterogeneity and 20.4 percent lower in the economy without frailty shocks.

Almost all of the difference between the variance of log lifetime earnings in the benchmark and the NFH economy is due to higher earnings at the bottom of the distribution in the NFH economy. Figure 7 displays the ratios of lifetime earnings at the 5th, 10th, 90th, and 95th percentile relative to the median by age in the the two economies. Notice that, after age 40, there are large differences across the two economies between the ratios of the 5th and 10th percentiles relative to the median. In contrast, there are little differences in the ratios of the 90th and 95th. Individuals in the bottom of the lifetime earnings distribution in the benchmark economy are more likely to be in poor health. They are also more likely to be less educated which means they face larger negative effects of poor health on their labor productivity. Giving these individuals the average frailty profile increases both their wages and their labor supply. In contrast, individuals at the top of the lifetime earnings distribution in the benchmark economy are mostly college-educated and healthy. As a result, giving these individuals the average frailty profile has little effect on their earnings.

Table 8: Variance of log lifetime earnings: the relative importance of frailty shocks versus initial frailty heterogeneity.

	Var. log lifetime earnings			
	age 45	age 55	age 65	age 75
Benchmark	0.384	0.438	0.437	0.405
No frailty heterogeneity	0.335	0.321	0.311	0.320
% change relative to benchmark	-12.9	-26.8	-28.9	-21.1
Removing only frailty shocks	0.343	0.349	0.349	0.369
% change relative to benchmark	-10.7	-20.4	-20.1	-8.8
Removing only frailty fixed effect	0.355	0.394	0.382	0.379
% change relative to benchmark	-7.7	-10.0	-12.5	-6.4

Note: In the "No frailty heterogeneity" counterfactual all individuals have the average frailty age profile. "Removing only frailty shocks" removes only ex post uncertainty/shocks but retains all the initial fixed-effect heterogeneity. "Removing only frailty fixed effect" only removes initial fixed effect heterogeneity but retains all the shocks and uncertainty.

6.1 Breaking down the effect of health inequality

Recall that there are five channels through which frailty can affect earnings inequality in the model: SSDI acceptance probabilities, labor productivity, disutility of work, amount of out-of-pocket medical expenditures, and mortality risk. How important are each of these channels for generating the differences in the variance of log lifetime earnings' profiles between the benchmark and NFH economies?

To assess the relative importance of each channel, we consider five additional counterfactual economies. Each economy is identical to the benchmark expect that, for one of the five channels, the impact of frailty is determined by the average frailty profile instead of a person's individual profile. Specifically, in counterfactual economy 1, labelled "NFH in SSDI", individuals' probability of successful SSDI application is determined by the average frailty profile. In counterfactual economy 2, labelled "NFH in Labor Productivity", individuals' labor productivity is determined by the average frailty profile. In counterfactual economy 3, "NFH in Disutility", disutility from working is determined by the average frailty profile. In counterfactual economy 4, "NFH in Medical Expenditure", out-of-pocket medical expenditures are determined by the average frailty profile. Finally, in counterfactual economy 5, "NFH in Mortality", mortality rates are determined by the average frailty profile.

The results of this decomposition exercise show that, according to the model, the SSDI program is the most important channel through which health inequality generates lifetime earnings inequality. Table 9 presents the differences in the variances of log lifetime earnings between the benchmark and each counterfactual economy at four ages. Notice that the labor productivity channel has the largest impact on lifetime earnings inequality at younger ages. Shutting down this channel reduces lifetime earnings inequality at age 45 by 5.6 percent, whereas shutting down the SSDI channel actually increases it by 5.1 percent. However, by age 55, the primary channel through which health inequality generates lifetime earning inequality is the SSDI channel. Shutting down the SSDI channel, reduces lifetime earnings inequality

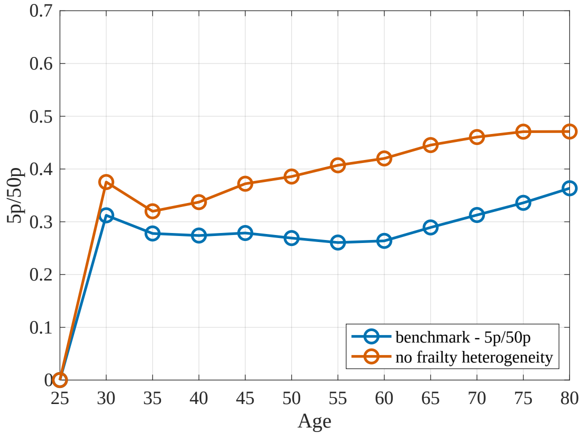

(a) 5th percentile to median

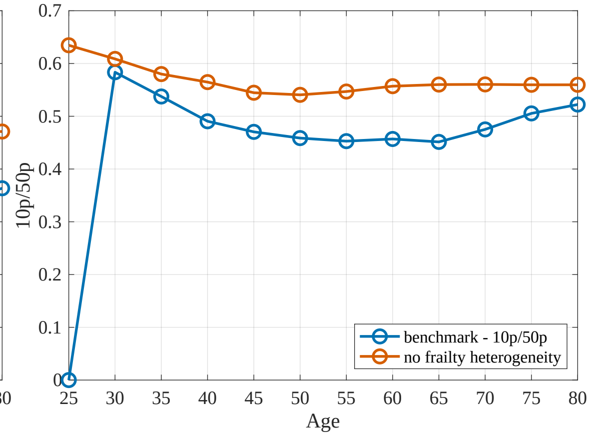

(b) 10th percentile to median

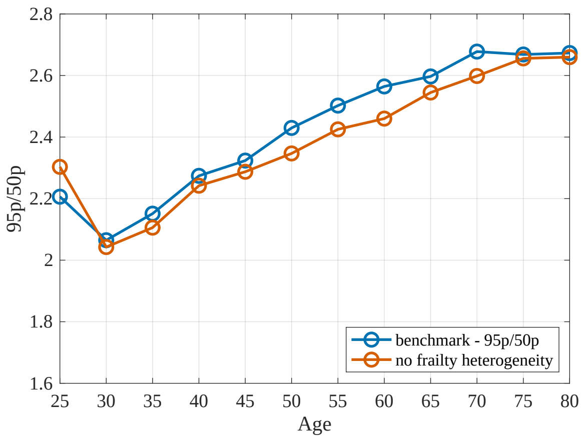

(c) 95th percentile to median

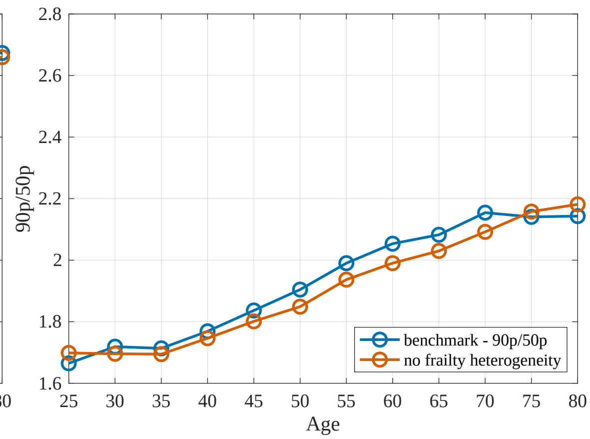

(d) 90th percentile to median

Figure 7: Inequality in lifetime earnings. Benchmark model (blue) vs model without heterogeneity in frailty (red).

by 8.1 percent at age 55. In contrast, shutting down the labor productivity channel, the second most important channel, reduces it by only 7.5 percent. At age 65, removing the SSDI channel reduces lifetime earnings inequality by 15.5 percent, nearly double that of the effect of removing the labor productivity channel. By age 75, the effect of removing the SSDI channel is triple that of removing the labor productivity channel.

Why does shutting down the SSDI channel actually increase lifetime earnings inequality at younger ages? In the benchmark economy, the SSDI program creates incentives for young frail people to work. These individuals have a high probability of getting SSDI transfers in the future and they want to accumulate earnings credit to raise their benefit in anticipation. Using the average frailty profile to determine SSDI eligibility substantially weakens this incentive as now the likelihood of a highly frail individual getting on SSDI is much lower. However, these individuals still suffer high disutility of work and have relatively low wages. These effects push young frail workers in the "NFH in SSDI" economy out of the labor force and onto means-tested programs. Figure 8 presents the labor force participation, DI

Table 9: Variance of log lifetime earnings: the relative importance of the different frailty channels.

	Var. log lifetime earnings (%Δ relative to benchmark)			
	age 45	age 55	age 65	age 75
NFH in SSDI	5.1	−8.1	−15.5	−14.9
NFH in Labor Prod.	−5.6	−7.5	−8.3	−4.9
NFH in Disutility	−1.6	−1.9	−2.3	−1.6
NFH in Med. Exp.	−0.4	−0.1	−0.3	−0.1
NFH in Mortality	−2.1	−1.0	6.9	7.0

Note: Each row shows the percentage change in the variance of log lifetime earnings in the counterfactual economy relative to the benchmark. Each counterfactual is identical to the benchmark expect that there is no frailty heterogeneity in the listed channel. Instead, the impact of frailty via that channel is determined by the average frailty age profile.

recipiency, and means-tested transfer recipiency rates for the top five percentiles of the frailty distribution in the benchmark and several of the counterfactual economies.[48] Consistent with the intuition above, the left panel of the figure shows that highly frail young people in the "NFH in SSDI" economy are less likely to be in the labor force than those in the benchmark, while the right panel shows that they are more likely to be on means-tested transfers. Thus, by reducing the incentives for young frail individuals to work, shutting down the SSDI channel, reduces labor force participation rates at younger ages and, hence, increases earnings inequality.

In contrast, shutting down the SSDI channel increases labor force participation rates of frail individuals at older ages as the first panel of Figure 8 shows. Compared to young highly frail individuals, older highly frail individuals are more likely to have worked and accumulated wealth. As a result, many are not eligible for means-tested programs. In the "NFH in SSDI" economy, given that they have a low probability of getting on SSDI, they continue to work. This impact of removing the SSDI channel on the labor supply of older frail individuals is the primary reason for the large and increasing decline in the variance of log lifetime earnings in the "NFH in SSDI" counterfactual relative to the benchmark.

The labor productivity channel is the second most important channel through which health inequality matters for lifetime earning inequality in the model. Shutting down this channel reduces lifetime earnings inequality at all ages. Using the average frailty profile to determine labor productivity reduces the variation in wages conditional on age. This has two effects. First, it leads directly to a reduction in earnings, and hence, lifetime earnings inequality. Second, it increases the returns from working for less educated individuals which increases their labor supply. This second effect operates even for the most highly frail individuals as Figure 8 shows.[49] Notice that, at all ages, these individuals have higher labor

[48] The effects of the SSDI channel on the labor-force participation, SSDI recipiency, and mean-tested transfer recipiency rates of the other frailty groups can be seen in Section 5 of the Online Appendix. The figures show that there is little impact of removing health inequality for individuals with frailty below the 70th percentile. For those in the 70th–95th percentiles, the effects of the SSDI channel are similar to those in Figure 8.

[49] See Section 5 of the Online Appendix for a summary of the overall effects of shutting down each channel

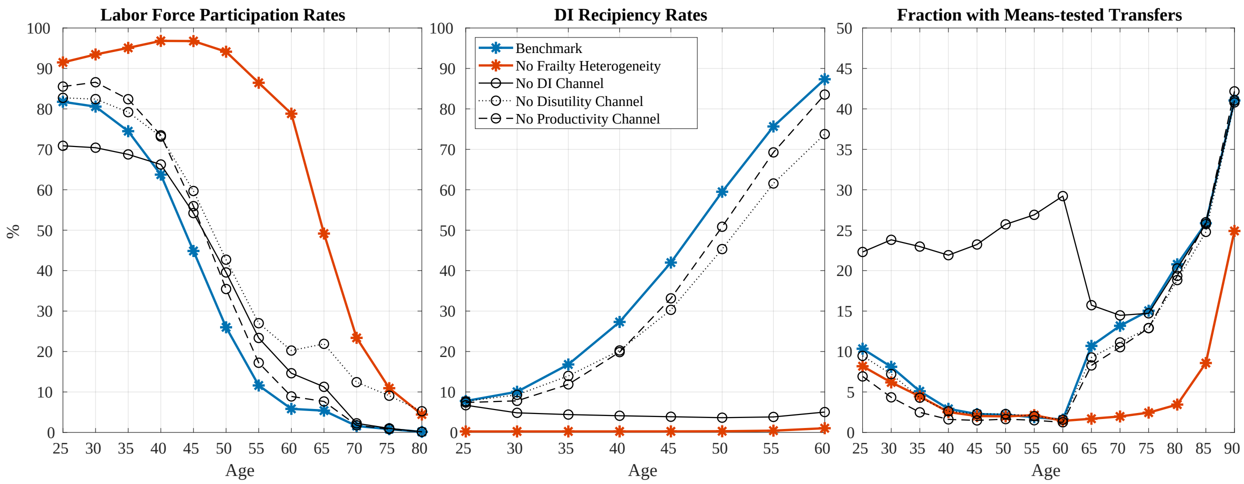

Figure 8: Labor force participation rates (left panel), DI recipiency rates (middle panel), and fraction receiving means-tested transfers (right panel) by age in the baseline economy, the no-frailty-heterogeneity economy, and counterfactuals 1–3.

force participation and lower DI and means-tested transfer recipiency rates in the "NFH in Labor Productivity" economy relative to the benchmark.

The disutility, medical expense, and mortality channels play relatively smaller roles. The smaller role of the mortality channel is due to two offsetting effects of shutting it down. First, it increases the life expectancy of frail individuals which increases their returns to work and labor supply. This effect works to reduce lifetime earning inequality. Second, since mortality and productivity are negatively correlated in the benchmark (due to both education and health), it raises the survival rates of individuals in the bottom of the lifetime earnings distribution relative to those in the top. This second effect, which grows with age due to the nature of mortality risk, works to increase lifetime earnings inequality.

6.2 Alternative measures of inequality

Our findings show that health inequality increases lifetime earnings inequality. The increase is driven by the negative impacts of poor health on the labor supply and earnings of individuals in the bottom of the lifetime earnings distribution. This increase is offset by SSDI and means-tested transfers. However, the offset is only partial. As a result, health inequality is also a significant driver of inequality in disposable income. Figure 9a presents the variance of log lifetime disposable income by age in the benchmark and NFH economies.[50] Comparing Figure 9a with Figure 6a shows that removing health inequality reduces lifetime disposable income inequality by roughly half as much as it reduces lifetime earnings inequality. For instance, the variance of log lifetime earnings at age 65 falls by 28.9 percent, while the vari-

on labor force, SSDI, and means-tested program participation rates.

[50]Disposable income as the sum of labor earnings and transfers net of all taxes. Lifetime disposable income at each age is the sum of disposable income to date. We have done all the calculations with an alternative definition that includes capital income. The results are very similar.

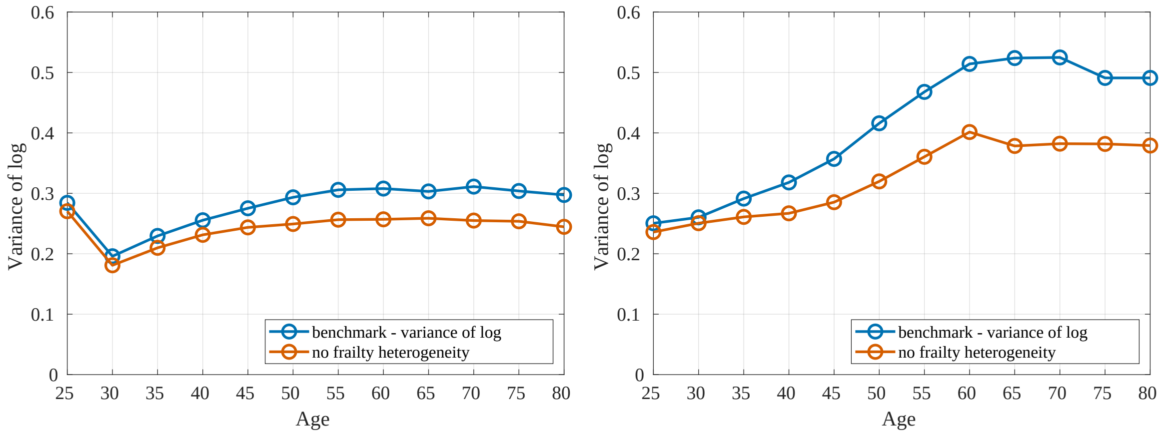

(a) Variance of log lifetime disposable income

(b) Variance of log consumption

Figure 9: Panel (a) is the variance of log lifetime disposable (defined as the sum of labor earnings and transfers net of taxes). Panel (b) is the variance of log consumption. The blue line is the benchmark and the red line is the economy with no frailty heterogeneity.

ance of log lifetime disposable income at age 65 falls by 14.7 percent.[51] This finding implies that about half of the inequality in lifetime earnings due to poor health is undone through the tax and transfer system.

Interestingly, the relatively smaller impacts of health inequality on lifetime disposable income inequality do not translate into smaller impacts of health inequality on consumption inequality. Figure 9b shows that removing health inequality reduces the variance of log current consumption by a similar amount as the variance of log lifetime earnings. The variance of log consumption at age 65 is 27.8 percent lower in the NFH economy as compared to the benchmark. Why are the impacts on consumption inequality so much larger than the impacts on lifetime disposable income inequality? The primary reason is that removing health inequality reduces wealth inequality. Less educated individuals in poor health in the benchmark have lower savings than less educated individuals in the NFH economy for three reasons. First, they have less income. Second, they have less incentives to accumulate wealth due to a lower life expectancy. Third, poor health (or the risk of poor health) means they are more impacted by the negative effects of the means-tested transfer program on incentives to save.[52]

6.3 Value of social security disability insurance

Our findings above indicate that health inequality is a major contributor to inequality in lifetime earnings. They also indicate that the primary channel through which health inequality generates lifetime earnings inequality is the SSDI program. The incentives for middle-aged

[51]See Section 5 of the Online Appendix for a counterpart to Table 8 for lifetime disposable income.

[52]It is well documented that means-tested transfer programs distort savings incentives and that their distortionary effects are larger for lower income individuals. See, for instance, Hubbard et al. (1995).

Table 10: Aggregate and welfare effects of eliminating SSDI

	Benchmark	Remove SSDI benefit & tax		
	(1)	(2)	(3)	(4)
Welfare – C.V. relative to benchmark (%)				
All	n.a.	-0.46	-0.84	-1.01
HSD	n.a.	-1.56	-1.90	-2.62
HSG	n.a.	-0.84	-1.21	-1.39
COL	n.a.	0.63	0.24	0.30
Inequality – Variance of				
log lifetime earnings (at age 65)	0.437	0.362	0.364	0.312
log lifetime disposable income (at age 65)	0.303	0.269	0.270	0.258
log consumption (overall)	0.173	0.155	0.156	0.138
Change relative to benchmark (%)				
GDP	n.a.	1.68	1.65	2.25
Consumption	n.a.	1.92	1.49	2.22
Capital	n.a.	1.68	1.65	2.25
Hours	n.a.	2.03	1.93	3.56
GDP per hour	n.a.	-0.34	-0.27	-1.27
Fraction of population (%)				
Working	87.65	89.27	89.17	90.70
On SSDI	4.58	0.00	0.00	0.00
On means tested transfers	4.77	6.41	6.51	4.98
Policy Variables				
Payroll tax rate (%)	12.40	11.42	11.42	11.42
Consumption floor (% of ave. earning)	11.00	11.00	11.00	10.46
Tax function parameter (λ)	0.91	0.91	0.90	0.91

Note: Column (2) is *partial equilibrium*:z removal of SSDI benefit and corresponding fraction of payroll tax. Column (3) and (4) are *general equilibrium* with government budget balanced by adjusting the income tax and consumption floor, respectively.

frail individuals to work and accumulate labor earnings are significantly reduced by the fact that they have a high probability of obtaining SSDI benefits if they apply. These results suggest that one way to reduce lifetime earnings inequality is to eliminate SSDI. We now assess the long-run welfare implications of such a policy.[53]

First, consider the direct impact of removing the SSDI program. This is done by setting the probability of getting SSDI to zero and adjusting the payroll tax so the total payroll tax receipt declines by exactly the amount of expenditures on SSDI benefits in the baseline.[54]

[53] We conduct this exercise in a model that lacks a private disability insurance market. We do this because the private disability insurance market in the U.S. is small and likely suffers from significant information frictions that impede its functioning. Only 3% of non-government workers have directly purchased this insurance and only 30% have obtained it indirectly through their employer. Hendren (2013) documents that coverage denial rates in the market are high and driven by the presence of private information.

[54] The initial distribution of individuals across employment states is also changed by moving DI beneficiary's

We refer to this experiment as *partial equilibrium*, since we do not attempt to balance the government budget. The second column of Table 10 shows the steady-state (long-run) welfare and aggregate implications. Notice that, consistent with the findings above, eliminating the SSDI program significantly reduces inequality in earnings, income and consumption. The variances of age-65 log lifetime earnings and disposable income fall by 17.2 and 11.2 percent, respectively. The variance of overall log consumption falls by 10.4 percent. It also increases aggregate consumption, as well as, aggregate GDP. Yet, it is ex ante welfare reducing. The decline in ex ante welfare, measured as the equivalent change in lifetime consumption, is 0.46 percent.

The decline in ex ante welfare when SSDI is removed is driven by the welfare losses of less educated individuals. Whereas college graduates benefit, on net, from this change (mostly due to lower taxes), both high school dropouts and high school graduates are worse off with welfare declines equivalent to 1.56 and 0.84 percent of lifetime consumption, respectively. For these less educated individuals, the SSDI program provides valuable insurance against the risk of becoming highly frail and incurring high disutility from work together with lower wages. Eliminating it leaves them more exposed to this risk. Their resulting utility losses dominate the ex ante welfare results.

Interestingly, removing the SSDI program has a differential impact on the labor supply of middle-aged versus younger frail workers. Middle-aged frail workers increase their labor supply as the cost of not working has gone up and they have already accumulated too much wealth to be eligible for means-tested transfers. Their response drives the rise in the aggregate labor force participation rate that can be seen in Table 10. The labor force participation rates of younger frail workers, under age 40, actually decline. This happens because removing the SSDI program has a similar effect on young frail individuals as using the average frailty profile to determine SSDI eligibility has on them. That is, it substantially reduces their incentive to work and accumulate lifetime earnings in anticipation of receiving SSDI transfers during middle-age. The loss of this work-incentive effect for young frail individuals is the primary driver of the rise in means-tested transfer recipiency rates reported in Table 10.

The rise in the means-tested transfer recipiency rates generates an increase in government expenditures. To show the impact of closing this fiscal gap on our welfare calculations, we perform two additional experiments. In the first experiment, we finance the expansion of means-tested programs by raising federal income taxes (by reducing parameter λ in the HSV tax function). The results of this experiment are reported in column (3) of Table 10. Alternatively, we close the gap by reducing the level of means-tested transfer benefits (by reducing consumption floor parameter $\bar{c}$). The results for this experiment are reported in column (4). These *general equilibrium* results reinforce our initial finding. Eliminating the SSDI program, despite reducing inequality in income and consumption and generating sizable gains in aggregate consumption and GDP, is not welfare improving. Welfare losses from eliminating the program are even larger in general equilibrium. If we close the fiscal gap by increasing income tax rates, ex ante welfare falls by 0.84 percent. If we close the gap instead by reducing the means-tested consumption floor, it falls by 1.01 percent. In both cases, welfare losses are particularly large for high school dropouts whose welfare declines by

to non-employed.

either 1.90 percent or 2.62 percent depending on which rule is used to clear the government budget.

Given our finding that the SSDI program is ex ante welfare improving, we now consider two additional sets of experiments that explore whether the current scale of the program is too low or too high. In the first set, we reduce SSDI benefits by 10 percent along with an appropriate adjustment in payroll taxes. We perform this experiment with and without balancing the government budget. The second set is the same as the first one except that we increase SSDI benefits by 10 percent.

The results, which are reported in Table 11, show that from an ex ante welfare perspective the current SSDI program may not be generous enough. First, a 10 percent decline in benefits leads to welfare losses which are sizable in comparison to the welfare changes reported in Table 10. For example, the overall ex ante welfare loss from cutting the program by 10 percent is 0.24 percent, more than a quarter of the 0.84 percent welfare loss from cutting the program all together. Second, even when the tax implications of balancing the government budget are accounted for, expanding benefits by 10 percent increases the welfare of high school dropouts and graduates by significantly more than the losses it generates for the college educated group. Overall, such a policy raises ex ante welfare by 0.20 percent.

Finally, notice in Table 11 that cutting DI benefits reduces the fraction of individuals receiving means-tested transfers, while raising benefits increases it. This may be surprising given that eliminating the DI program all together increases means-tested transfer recipiency rates. Recall that the primary reason that recipiency rates rise when the SSDI program is removed is because frail individuals under age 45 have a reduced incentive to work and accumulate SSDI credits. While eliminating SSDI increases the means-tested transfer recipiency rates of this group, cutting SSDI benefits by 10 percent decreases them. The reason is that with lower (but still positive) SSDI benefit levels, these individuals have a reduced incentive to exit the labor force and collect MTSI transfers while trying to get on SSDI. Consequently, they have higher labor force participation rates and lower means-tested transfer recipiency rates both when young and during retirement. This same effect, operating in reverse, is also the primary reason that raising SSDI benefits by 10 percent raises means-tested transfer recipiency rates. Higher SSDI benefits increase the incentives for frail individuals to exit the labor force at relatively younger ages, increasing their likelihood of means-tested transfer eligibility both when young and after retirement. These findings highlight how the subtle interactions between these two programs make it difficult to predict how changes to one will impact the size of the other.[55]

7 Conclusion

In this paper, we document empirically that declines in health reduce labor productivity and the probability of labor force participation. The effects are concentrated in less educated individuals and those already in poor health suggesting that health inequality may be an important source of lifetime earnings inequality. Using a structural model, we demonstrate that this is indeed the case: almost 29% of the variation in lifetime earnings at age 65 is due

[55] Low and Pistaferri (2015) document similar non-monotonic effects of increasing the generosity of means-tested food stamps on DI application rates.

Table 11: Aggregate and welfare effects of marginal changes to SSDI

	Benchmark	Cut SSDI benefit by 10%		Raise SSDI benefit by 10%	
	(1)	(2)	(3)	(4)	(5)
Welfare – C.V. relative to benchmark (%)					
All	n.a.	-0.31	-0.24	0.28	0.20
HSD	n.a.	-0.87	-0.80	0.79	0.72
HSG	n.a.	-0.40	-0.32	0.36	0.28
COL	n.a.	0.06	0.14	-0.06	-0.15
Change relative to benchmark (%)					
GDP	n.a.	0.24	0.25	-0.23	-0.23
Consumption	n.a.	0.22	0.31	-0.20	-0.31
Capital	n.a.	0.24	0.25	-0.23	-0.23
Labor input	n.a.	0.24	0.25	-0.23	-0.23
Hours	n.a.	0.49	0.51	-0.44	-0.46
GDP per hour	n.a.	-0.24	-0.25	0.21	0.23
Fraction of population (%)					
Working	87.65	88.12	88.14	87.23	87.20
On SSDI	4.58	4.31	4.31	4.83	4.83
On means tested transfers	4.77	4.73	4.72	4.91	4.93
Policy Variables					
Payroll tax rate (%)	12.40	12.25	12.25	12.55	12.55
Consumption floor (% of ave. earning)	11.00	11.00	11.00	11.00	11.00
Tax function parameter (λ)	0.91	0.91	0.91	0.91	0.91

Note: Column (1) is the benchmark. Column (2) and (4) are *partial equilibrium*, meaning only payroll tax adjusts by the amount of change in SSDI benefit, but consolidated government budget is not balanced. Column (3) and (5) are *general equilibrium*, meaning government consolidated budget is balanced by changing income tax.

to the fact that individuals in the United States face risky and heterogeneous lifecycle health profiles. A decomposition exercise shows that the impact of poor health on access to social security disability benefits is the most important factor driving our results. In other words, the primary reason why individuals in poor health have low lifetime earnings is because they have a high likelihood of obtaining social security disability benefits. The negative effects of poor health on the wages of less educated workers also play an important role. Interestingly, we find that the disutility effects of working while in poor health play a relatively small role. These findings indicate that the social security program is an important contributor to lifetime earnings inequality. Despite this, we document that it is ex ante welfare improving and, if anything, should be expanded.

A Appendix

A.1 Reverse causality: estimated effects of earnings on health

The empirical analysis in Section 2 focuses on estimating the causal effect of frailty on earnings, hours, and wages. However, it is also possible that earnings affects frailty. To examine this proposition we estimate the following regression

$$f_{i,t} \quad = \quad b_i + \gamma y_{i,t} + \alpha_1 f_{i,t-1} + \alpha_2 f_{i,t-2} + \beta \boldsymbol{Z}_{i,t} + \varepsilon_{i,t}, \tag{19}$$

using the system GMM estimation procedure outlined in Section 2.2. Here, $f_{i,t}$ is the level of frailty of individual i at date t and $y_{i,t}$ is their date t log earnings. The set of controls $\boldsymbol{Z}_{i,t}$ is the same as in those in Section 2. We run the same set of regressions as we did with log earnings on the left-hand-side. I.e., we look at the effect of log earnings on frailty overall and by education, health, and age. To get valid instruments for these regressions, we must go back further in lag length then we did for the regressions in Section 2. We also need to include more instruments. We use the sixth through eighth lags of frailty and log earnings as instruments.

The results are presented in Table 12. Notice that we find no evidence that earnings affects frailty. In all cases, the effects of log earnings on frailty are small and in only one regression do we see significance at the 10 percent level. Note also that, in all specifications, the diagnostic tests are passed. In other words, the error terms are not second order autoregressive and the Hansen-Sargan tests indicate that the null hypothesis of valid instruments cannot be rejected.

Table 12: Effect of Earnings on Frailty

	Everyone				Workers			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
frailty_{t-1}	0.445	0.334	-0.152	-0.456	-0.182	0.712*	0.302	-0.190
	(0.463)	(0.435)	(0.528)	(0.400)	(0.566)	(0.416)	(0.737)	(0.498)
frailty_{t-2}	0.602	0.661	1.124**	1.446***	1.316**	0.405	0.820	1.321**
	(0.447)	(0.443)	(0.495)	(0.404)	(0.596)	(0.451)	(0.741)	(0.529)
$\log(\text{earnings}_t)$	0.004*				-0.004			
	(0.002)				(0.007)			
$\log(\text{earnings}_t)$ × HSD		0.003				0.001		
		(0.002)				(0.002)		
$\log(\text{earnings}_t)$ × HS		-0.008				-0.019		
		(0.039)				(0.073)		
$\log(\text{earnings}_t)$ × CL		0.000				-0.000		
		(0.001)				(0.001)		
$\log(\text{earnings}_t)$ × Good Health			0.000				0.000	
			(0.003)				(0.008)	
$\log(\text{earnings}_t)$ × Bad Health			0.002				0.000	
			(0.002)				(0.008)	
$\log(\text{earnings}_t)$ × Young				-0.000				-0.004
				(0.001)				(0.007)
$\log(\text{earnings}_t)$ × Old				-0.000				-0.003
				(0.002)				(0.007)
Controls	YES	YES	YES	YES	YES	YES	YES	YES
Observations	50,844	50,844	50,844	50,844	27,444	27,444	27,444	27,444
AR(1) test (p-value)	0.531	0.573	0.501	0.001	0.260	0.388	0.763	0.188
AR(2) test (p-value)	0.333	0.260	0.061	0.002	0.060	0.570	0.380	0.032
Hansen test (p-value)	0.269	0.842	0.621	0.129	0.440	0.430	0.747	0.848
Diff-in-Hansen test (p-value)	0.450	0.852	0.894	0.132	0.656	0.225	0.805	0.818
Diff-in-Hansen test (p-value), Y-lag set	.	0.990	0.723	0.223	.	0.245	0.814	0.788

Notes: Columns (1)–(4) show regression results for the entire sample, regardless of employment status. Columns (5)–(8) show results conditional on continued employment. All regressions include controls (marital status, marital status interacted with gender, number of kids, number of kids interacted with gender, year dummies, and a quadratic in age). ‘HSD’ is high school dropout, ‘HS’ is high school graduate, and ‘CL’ is college graduate. ‘Good/Bad Health’ is frailty below/above the 75th percentile. ’Young/Old’ are individuals younger/older than 45 years of age. Standard errors are in parenthesis. *$p < 0.1$; **$p < 0.05$; ***$p < 0.01$.

References

Aizawa, N., S. Kim and S. Rhee, "Labor Market Screening and Social Insurance Program Design for the Disabled," (2020). 5

Al-Sadoon, M. M., S. Jiménez-Martín, J. M. Labeaga et al., "Simple methods for consistent estimation of dynamic panel data sample selection models," Technical Report, 2019. 3, 23

Arellano, M. and S. Bond, "Some Tests of Specification for Panel Data: Monte Carol Evidence and and Application to Employment Equations," *The Review of Economic Studies* 58 (1991), 277–297. 9

Arellano, M. and O. Bover, "Another look at the instrumental variable estimation of error-components models," *Journal of econometrics* 68 (1995), 29–51. 10

Attanasio, O. P., "Consumption," *Handbook of macroeconomics* 1 (1999), 741–812. 22

Autor, D. H. and M. G. Duggan, "The rise in the disability rolls and the decline in unemployment," *The Quarterly Journal of Economics* 118 (2003), 157–206. 5

Beck, T., R. Levine and N. Loayza, "Finance and the Sources of Growth," *Journal of financial economics* 58 (2000), 261–300. 9, 10

Bell, F. C. and M. L. Miller, "Life tables for the United States social security area," *Social Security Administration Publications* (2005). 21

Benabou, R., "Tax and Education Policy in a Heterogeneous Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency?," *Econometrica* 70(2) (2002), 481–517. 27

Blundell, R. and S. Bond, "Initial conditions and moment restrictions in dynamic panel data models," *Journal of econometrics* 87 (1998), 115–143. 3, 9, 10

———, "GMM estimation with persistent panel data: an application to production functions," *Econometric reviews* 19 (2000), 321–340. 10

Blundell, R., J. Britton, D. MC and E. French, "The impact of health on labor supply near retirement," mimeo, 2017. 5

Blundell, R. and T. MaCurdy, "Labor supply: A review of alternative approaches," in *Handbook of labor economics*volume 3 (Elsevier, 1999), 1559–1695. 22

Bond, S. and C. Meghir, "Dynamic investment models and the firm's financial policy," *The Review of Economic Studies* 61 (1994), 197–222. 9

Bond, S. R., "Dynamic panel data models: a guide to micro data methods and practice," *Portuguese economic journal* 1 (2002), 141–162. 10

Bound, J., M. Schoenbaum, T. R. Stinebrickner and T. Waidmann, "The dynamic effects of health on the labor force transitions of older workers," *Labour Economics* 6 (1999), 179–202. 5

Braun, R. A., K. Kopecky and T. Koreshkova, "Old, Sick, Alone and Poor: A Welfare Analysis of Old-Age Social Insurance Programs," Forthcoming Review of Economic Studies, 2015. 5

Capatina, E., "Life-cycle effects of health risk," *Journal of Monetary Economics* 74 (2015), 67–88. 5

Caselli, F., G. Esquivel and F. Lefort, "Reopening the convergence debate: a new look at cross-country growth empirics," *Journal of economic growth* 1 (1996), 363–389. 9

Dalgaard, C.-J. and H. Strulik, "Optimal Aging and Death: Understanding the Preston Curve," *Journal of the European Economic Association* 12 (2014), 672–701. 2

De Nardi, M., E. French and J. B. Jones, "Why do the Elderly Save? The Role of Medical Expenses," *Journal of Political Economy* 118 (2010), 37–75. 5

De Nardi, M., S. Pashchenko and P. Porapakkarm, "The Lifetime Costs of Bad Health," Working Paper, 2017. 5

Dobkin, C., A. Finkelstein, R. Kluender and M. J. Notowidigdo, "The economic consequences of hospital admissions," *American Economic Review* 108 (2018), 308–52. 5

French, E., "The effects of health, wealth, and wages on labour supply and retirement behaviour," *The Review of Economic Studies* 72 (2005), 395–427. 5

French, E. and J. Song, "The effect of disability insurance receipt on labor supply," *American economic Journal: economic policy* 6 (2014), 291–337. 25, 26, 29

García-Gómez, P., H. Van Kippersluis, O. O'Donnell and E. Van Doorslaer, "Long-term and spillover effects of health shocks on employment and income," *Journal of Human Resources* 48 (2013), 873–909. 5

Goggins, W. B., J. Woo, A. Sham and S. C. Ho, "Frailty index as a measure of biological age in a Chinese population," *The Journals of Gerontology Series A: Biological Sciences and Medical Sciences* 60 (2005), 1046–1051. 6

Gomme, P. and P. Rupert, "Theory, Measurement and Calibration of Macroeconomic Models," *Journal of Monetary Economics* 54 (2007), 460–497. 28

Guner, N., R. Kaygusuz and G. Ventura, "Income taxation of US households: Facts and parametric estimates," *Review of Economic Dynamics* 17 (2014), 559–581. 28

Guvenen, F., "An Empirical Investigation of Labor Income Processes," *Review of Economic Dynamics* 12 (2009), 58–79. 24

Guvenen, F., G. Kaplan, J. Song and J. Weidner, "Lifetime incomes in the United States over six decades," Technical Report, National Bureau of Economic Research, 2017. 30

Heathcote, J., K. Storesletten and G. Violante, "Optimal Tax Progressivity: An Analytical Framework," *Quarterly Journal of Economics* 132(4) (2017), 1693–1754. 27

Heinesen, E. and C. Kolodziejczyk, "Effects of breast and colorectal cancer on labour market outcomes-average effects and educational gradients," *Journal of health economics* 32 (2013), 1028–1042. 5

Hendren, N., "Private Information and Insurance Rejections," *Econometrica* 81 (2013), 1713–1762. 38

Hoechle, D., M. Schmid, I. Walter and D. Yermack, "How much of the diversification discount can be explained by poor corporate governance?," *Journal of Financial Economics* 103 (2012), 41–60. 9

Holtz-Eakin, D., W. Newey and H. S. Rosen, "Estimating vector autoregressions with panel data," *Econometrica: Journal of the Econometric Society* (1988), 1371–1395. 9

Hong, J. H. and J.-V. Ríos-Rull, "Life insurance and household consumption," *American Economic Review* 102 (2012), 3701–30. 28

Hosseini, R., K. Kopecky and K. Zhao, "The Evolution of Health over the Life Cycle," Federal Reserve Bank of Atlanta Working Paper, 2019. 2, 6, 24

Hubbard, R. G., J. Skinner and S. P. Zeldes, "Precautionary Savings and Social Insurance," *Journal of Political Economy* 103 (1995), 360–399. 37

Imrohoroglu, A. and K. Zhao, "The Chinese Saving Rate: Long-Term Care Risks, Family Insurance, and Demographics," *Journal of Monetary Economics* 96 (2018). 5

Jeon, S.-H., "The long-term effects of cancer on employment and earnings," *Health economics* 26 (2017), 671–684. 5

Jones, J. B. and Y. Li, "The effects of collecting income taxes on social security benefits," *Journal of Public Economics* 159 (2018), 128–145. 18

Kim, S. and S. Rhee, "Understanding the aggregate effects of disability insurance," (2020). 5

Kitao, S., "A life-cycle model of unemployment and disability insurance," *Journal of Monetary Economics* 68 (2014), 1–18. 5

Kopecky, K. and T. Koreshkova, "The Impact of Medical and Nursing Home Expenses on Savings," *American Economic Journal: Macroeconomics* 6 (2014), 29–72. 5, 28

Kulminski, A., A. Yashin, K. Arbeev, I. Akushevich, S. Ukraintseva, K. Land and K. Manton, "Cumulative index of health disorders as an indicator of aging-associated processes in the elderly: results from analyses of the National Long Term Care Survey," *Mechanisms of ageing and development* 128 (2007a), 250–258. 6

Kulminski, A. M., S. V. Ukraintseva, I. V. Akushevich, K. G. Arbeev and A. I. Yashin, "Cumulative index of health deficiencies as a characteristic of long life," *Journal of the American Geriatrics Society* 55 (2007b), 935–940. 6

Low, H. and L. Pistaferri, "Disability insurance and the dynamics of the incentive insurance trade-off," *American Economic Review* 105 (2015), 2986–3029. 5, 23, 40

Lundborg, P., M. Nilsson and J. Vikström, "Heterogeneity in the impact of health shocks on labour outcomes: evidence from Swedish workers," *Oxford Economic Papers* 67 (2015), 715–739. 5

Maestas, N., K. J. Mullen and A. Strand, "Does disability insurance receipt discourage work? Using examiner assignment to estimate causal effects of SSDI receipt," *American economic review* 103 (2013), 1797–1829. 5, 15

Meyer, B. D. and W. K. Mok, "Disability, earnings, income and consumption," *Journal of Public Economics* 171 (2019), 51–69. 5

Michaud, A. and D. Wiczer, "The Disability Option: Labor Market Dynamics with Macroeconomic and Health Risks," (2017). 5

Mitnitski, A., X. Song, I. Skoog, G. A. Broe, J. L. Cox, E. Grunfeld and K. Rockwood, "Relative fitness and frailty of elderly men and women in developed countries and their relationship with mortality," *Journal of the American Geriatrics Society* 53 (2005), 2184–2189. 6

Mitnitski, A. B., J. E. Graham, A. J. Mogilner and K. Rockwood, "Frailty, fitness and late-life mortality in relation to chronological and biological age," *BMC geriatrics* 2 (2002), 1. 6

Mitnitski, A. B., A. J. Mogilner and K. Rockwood, "Accumulation of deficits as a proxy measure of aging," *The Scientific World Journal* 1 (2001), 323–336. 6

Mitnitski, A. B., X. Song and K. Rockwood, "The estimation of relative fitness and frailty in community-dwelling older adults using self-report data," *The Journals of Gerontology Series A: Biological Sciences and Medical Sciences* 59 (2004), M627–M632. 6

Nickell, S., "Biases in dynamic models with fixed effects," *Econometrica: Journal of the Econometric Society* (1981), 1417–1426. 9

Ozcan, S., "Preventive vs. Curative Medicine: A Macroeconomic Analysis of Health Care over the Life Cycle," , 2013. 5

Pohl, V., C. Neilson and F. Parro, "The Effect of Health Shocks on Employment: Evidence From Accidents in Chile," (2013). 5

Prados, M. J., "Health And Earnings Inequality Over The Life Cycle And The Redistributive Potential Of Health Policies," University of Southern California Working Paper, 2017. 5

Rockwood, K., M. Andrew and A. Mitnitski, "A comparison of two approaches to measuring frailty in elderly people," *The Journals of Gerontology Series A: Biological Sciences and Medical Sciences* 62 (2007), 738–743. 6

Rockwood, K. and A. Mitnitski, "Frailty in relation to the accumulation of deficits," *The Journals of Gerontology Series A: Biological Sciences and Medical Sciences* 62 (2007), 722–727. 6

Roodman, D., "A note on the theme of too many instruments," *Oxford Bulletin of Economics and statistics* 71 (2009), 135–158. 10

Schünemann, J., H. Strulik and T. Trimborn, "The gender gap in mortality: How much is explained by behavior?," *Journal of Health Economics* 54 (2017a), 79–90. 2

———, "Going from bad to worse: Adaptation to poor health spending, longevity, and the value of life," *Journal of Economic Behavior & Organization* 140 (2017b), 130–146. 2

Searle, S., A. Mitnitski, E. Gahbauer, T. Gill and K. Rockwood, "A standard procedure for creating a frailty index," *BMC Geriatrics* 8 (2008), 1. 6

Suen, R., "Technological Advance and the Growth in Health Care Spending," , 2006. 5

Whited, T. M., "Investment and financial asset accumulation," *Journal of Financial Intermediation* 1 (1991), 307–334. 9

Wintoki, M. B., J. S. Linck and J. M. Netter, "Endogeneity and the dynamics of internal corporate governance," *Journal of Financial Economics* 105 (2012), 581–606. 9

Woo, J., W. Goggins, A. Sham and S. Ho, "Social determinants of frailty," *Gerontology* 51 (2005), 402–408. 6

Wooldridge, J. M., *Econometric analysis of cross section and panel data* (MIT press, 2010). 9

Zhao, K., "Social Security and the Rise in Health Spending," *Journal of Monetary Economics* 64 (2014), 21–37. 5