

# Health-Dependent Preferences, Consumption, and Insurance

Nicolo Russo\*  
Goethe University Frankfurt

February 20, 2026

## Abstract

How does health affect preferences, self-insurance, and the value of government insurance? I build a life-cycle model in which health affects survival, earnings, medical expenses, and the marginal utility of consumption, and calibrate it using data from the Panel Study of Income Dynamics. My main findings are as follows. Bad health reduces the marginal utility of consumption, lowering savings over the life cycle and consumption in old age. A model without health-dependent preferences does not replicate the degree of self-insurance against health shocks observed in the data. Finally, health-dependent preferences reduce the household valuation of means-tested government insurance programs.

---

\*Goethe University Frankfurt, russo@econ.uni-frankfurt.de. I would like to thank Mariacristina De Nardi for her invaluable guidance and support. I would also like to thank Filip Babalievsky, Margherita Borella, Lidia Cruces, Dhananjay Ghei, Martin Garcia Vazquez, Alessandra Fogli, Jeremy Lise, Jo Mullins, Rory McGee, Fang Yang, and all the participants of the Macro-Micro workshop at the University of Minnesota for their advice and suggestions. Financial support from the German Research Foundation, grant no. 462655750 is gratefully acknowledged.

# I Introduction

Health shapes economic well-being. Bad health shortens life, lowers earnings, and raises medical expenses, but may also affect the marginal utility of consumption. While the first set of effects is well understood, there is limited evidence and no consensus on how health affects the marginal utility of consumption. This paper studies how health affects household preferences, consumption, saving, self-insurance, and the valuation of means-tested transfers in a quantitative life-cycle model that incorporates both channels.

In standard models without preference shocks, households choose consumption to smooth marginal utility across states. If health does not affect preferences, smoothing marginal utility is equivalent to smoothing consumption itself. If, instead, as suggested by previous literature, bad health affects the marginal utility of consumption, optimal consumption varies with health. Then, fluctuations in health affect consumption and saving behavior, the effectiveness of self-insurance, and the value of means-tested government transfers.

Measuring health is essential to studying these mechanisms. To do so, I use data from the Panel Study of Income Dynamics (PSID) and measure health with the frailty index. The frailty index originates from the medical literature and combines information on diseases, difficulties in activities of daily living, harmful lifestyle habits, and other health indicators into a single measure. In particular, frailty is defined as the share of health conditions out of all the possible ones.

I build and calibrate a life-cycle model of consumption and saving in which households, starting at age 25, face uncertainty about health, medical expenses, and longevity. Health and earnings each include persistent and transitory components. After retirement, households receive Social Security benefits.

In the model, health affects survival, earnings, medical expenses, and the marginal utility of consumption. As in De Nardi, French, and Jones (2010) and Keane, Capatina, and Maruyama (2020), medical expenses are exogenous and affect household resources. Given realistic uncertainty and government insurance, households optimally choose consumption and saving, while the government provides means-tested transfers that ensure a minimum consumption level.

I calibrate the model in two steps using PSID data. First, I estimate all components that can be identified outside the model, such as the processes for earnings and health. Second, I calibrate the government-provided consumption floor and the effect of health on the marginal utility of consumption by matching the consumption response to transitory earnings and health shocks.

My identification strategy, one of the paper’s main contributions, hinges on how consumption responds to transitory shocks. Consumption fluctuations reveal information about both the means-tested consumption floor and the effect of health on the marginal utility of consumption.

A transitory earnings shock affects current resources. If the consumption floor is high, more households consume at that floor, and their consumption responds little to the earnings shock. If the floor is low, fewer households are constrained, and their consumption responds more to the shock. Hence, the average consumption response to a transitory earnings shock identifies the consumption floor.

A transitory health shock affects consumption through two channels. It changes resources by affecting income and medical expenses (*resource channel*), and it changes the marginal utility of consumption directly (*marginal utility channel*). My model explicitly acknowledges the resource channel. If health had no effect on marginal utility, the consumption response to a health shock would be, for a given resource effect size, the same as that of an earnings shock. Once the resource channel is identified, the marginal utility channel is identified residually from the consumption response to a transitory health shock. The model also accounts for the small effect of transitory health shocks on life expectancy, which is therefore “netted out” of the consumption response to a transitory health shock, ensuring the validity of my identification strategy.

Persistent earnings and health shocks have additional effects compared to their transitory counterparts. A persistent earnings shock also affects the distribution of future earnings. A persistent health shock has even broader consequences because it impacts future health and, thus, life expectancy, future medical expenses, and future earnings. My model is designed to capture these effects.

Estimating the consumption response to persistent shocks in the data, however, requires more stringent assumptions than those for transitory shocks. In particular, identifying the consumption response to transitory shocks requires only that the laws of motion of the underlying processes be well specified and that consumption be independent of future shocks. In contrast, identifying the consumption response to persistent shocks requires assuming that log consumption evolves as a random walk, an assumption rejected both in the data and in life-cycle models with precautionary savings (Carroll (1997) and Commault (2022)). Moreover, it typically assumes permanent rather than persistent shocks, or at least that the persistence of the shocks is known in advance (see Kaplan and Violante (2010) on this point). For these reasons, I examine my model's implications for consumption responses to persistent shocks but do not use them for identification.

The paper delivers four main findings. First, bad health lowers the marginal utility of consumption. This finding is consistent with, among others, Finkelstein, Luttmer, and Notowidigdo (2013). The calibration also yields a consumption floor of about \$2,600 per person per year, which is within the range of values from previous literature.

Second, health-dependent preferences affect consumption and savings decisions over the life cycle. Most studies on the determinants of savings abstract from the effect of health on preferences. However, comparing the predictions of my baseline model with those of one in which health does not affect the marginal utility of consumption shows that, for instance, households in their eighties would consume 13% more and save 30% more if health had no effect on preferences. This is because health worsens over the life cycle, and the marginal utility of consumption is lower in worse health. As households desire to consume less when in worse health and health deteriorates with age, savings are also lower.

Third, health-dependent preferences affect the degree of self-insurance against health shocks. The standard approach to estimating self-insurance involves comparing changes in income and consumption (Blundell, Pistaferri, and Preston, 2008). However, it typically ignores how health affects both income and preferences (except for Blundell et al. (2024), which focuses on people aged 65 and older and does not calibrate a structural model). I show that a model without health-dependent preferences overstates the degree of self-insurance against health shocks compared to the data. In particular, the consumption response to a

transitory health shock predicted by a model without health-dependent preferences is about 40% smaller than the data estimate.

Fourth, health-dependent preferences affect households' valuation of means-tested government insurance. This is the first quantitative model to assess this channel. I show that health-dependent preferences reduce the household valuation of government insurance because, when a bad health shock hits, the marginal utility of consumption is lower with health-dependent preferences, and households need a lower consumption floor. I also show that government insurance is more valuable for low-income and less healthy households. Households in the bottom 5% of the earnings distribution value means-tested government insurance almost seven times as much as the population average, while those in the top 5% of the earnings distribution do not value it at all. Similarly, the welfare effect of a 50% reduction in means-tested programs is 24% larger for the sickest households (those in the top 5% of the frailty distribution) than for the healthiest ones (those in the bottom 5% of the frailty distribution).

The remainder of the paper proceeds as follows. Section II discusses the relationship with the literature and the main contributions. Section III presents the quantitative model. Section IV describes the data and the measurement of health. Section V outlines the empirical strategy and presents the calibration results. Section VI shows the effects of health-dependent preferences on consumption and savings, Section VII analyzes the effects of health-dependent preferences on self-insurance against health shocks, and Section VIII discusses the welfare effects of reforming means-tested government insurance. Section IX concludes.

## II Relationship to the literature and contributions

My paper relates to three strands of the literature and contributes to each of them. First, my paper relates to the literature on **health-dependent preferences**, which has not reached a consensus on the sign or magnitude of the effect of health on the marginal utility of consumption, in part because studies rely on different sources of variation and identification strategies. Two approaches have emerged in this literature. The first one is empirical and includes, among others, Viscusi and Evans (1990), Evans and Viscusi (1991), Sloan et al.

(1998), Edwards (2008), Finkelstein, Luttmer, and Notowidigdo (2013), Brown, Goda, and McGarry (2016), Gyrd-Hansen (2017), Kools and Knoef (2019), and Achou et al. (2023). Findings differ across studies: Evans and Viscusi (1991) finds no evidence of an effect of health on preferences; Finkelstein, Luttmer, and Notowidigdo (2013) shows that bad health reduces the marginal utility of non-medical consumption; Kools and Knoef (2019) finds the opposite. The second approach uses structural life-cycle models that allow health to shift the marginal utility of consumption and estimate this channel within the model. Papers following this approach include Lillard and Weiss (1997), Rust and Phelan (1997), De Nardi, French, and Jones (2010), Hong, Pijoan-Mas, and Ríos-Rull (2015), Koijen, Van Nieuwerburgh, and Yogo (2016), and Ameriks et al. (2020). Results also differ: Lillard and Weiss (1997) finds that marginal utility increases as health deteriorates; De Nardi, French, and Jones (2010) finds that it decreases; Hong, Pijoan-Mas, and Ríos-Rull (2015) finds that the effect is negative at age 65 and positive at older ages.<sup>1</sup> I contribute to this literature by providing a new way to identify the effect of health on preferences. I also provide the first quantitative assessment of the consequences of health-dependent preferences on self-insurance and government insurance.<sup>2</sup>

Second, my paper relates to the literature on **consumption insurance**. Notable papers in this literature are Cochrane (1991), Attanasio and Davis (1996), Blundell, Pistaferri, and Preston (2008), Kaplan and Violante (2010), Blundell, Pistaferri, and Saporta-Eksten (2016), Blundell, Pistaferri, and Saporta-Eksten (2018), Wu and Krueger (2018), Commault (2022), and Blundell et al. (2024). Numerous studies in this literature—including Blundell, Pistaferri, and Preston (2008)—focus solely on income risk during the working age. More recently, Blundell et al. (2024) considers income and health risks but focuses on the elderly

---

<sup>1</sup>My paper is closely related to Ameriks et al. (2020), as both build and estimate models featuring health-dependent preferences. They find that the marginal utility of consumption while in need of long-term care (LTC) is higher than when in better health states. Our models differ along many dimensions, including that they use discrete health (while I use continuous) and allow health to affect marginal utility only when in need of LTC (while I allow the effect to vary with frailty throughout the health distribution). Moreover, they identify the effect of LTC needs on utility using strategic survey questions, which differ substantially from the information I have access to in the PSID. These differences help explain why we find different results.

<sup>2</sup>Finkelstein, Luttmer, and Notowidigdo (2013) uses a stylized two-period model to evaluate the consequences of health-dependent preferences on savings for retirement, but does not evaluate the welfare effects of health-dependent preferences. I use a richer model and evaluate the consequences of health-dependent preferences on insurance.

and does not estimate a structural model. I contribute to this literature by building and calibrating a life-cycle model that incorporates earnings and health risk across the entire life cycle, and by clarifying how health-dependent preferences shape the interpretation of consumption responses to health shocks.

Third, my paper relates to the literature on **structural life-cycle models with health risk**. Seminal contributions in this literature are Palumbo (1999), French (2005), and De Nardi, French, and Jones (2010). This literature has expanded substantially and includes, among others, Scholz and Seshadri (2011), Kopecky and Koreshkova (2014), Capatina (2015), Braun, Kopecky, and Koreshkova (2016), De Nardi, French, and Jones (2016), De Nardi, Pashchenko, and Porapakkarm (2017), and recent contributions by Hosseini, Kopecky, and Zhao (2020), Keane, Capatina, and Maruyama (2020), Salvati (2020), Bolt (2021), Achou (2021), Dal Bianco and Moro (2022), Dal Bianco (2023), and Michaud and St. Amour (2023). In most of these models, health affects survival, earnings, and medical expenditures, but not the marginal utility of consumption directly. I extend this framework by allowing health to enter preferences and by quantifying how this channel alters saving behavior and the valuation of insurance programs.

### III Model

Households enter the model at age 25, retire exogenously at age 63, and die with certainty by age 89. They face health, earnings, and survival risk until retirement; after retirement, earnings risk is resolved. Each model period lasts two years, consistent with the biennial structure of the PSID.

Households can only invest in a risk-free asset with a constant rate of return. There are no annuity markets to insure against mortality risk, and accidental bequests are lost to the economy.

**Health-Dependent Preferences.** In each period, utility depends on the consumption of non-durable and non-medical goods,  $c_t$ , and frailty,  $f_t$ . The period flow utility of consumption

is:

$$u(c_t, f_t) = \delta(f_t) \frac{c_t^{1-\gamma}}{1-\gamma}. \quad (1)$$

The parameter  $\gamma$  is the coefficient of relative risk aversion, and  $\delta(f_t)$  measures the effect of health on marginal utility. Following Palumbo (1999) and De Nardi, French, and Jones (2010), I model the effect of health on marginal utility as:

$$\delta(f_t) = 1 + \delta f_t. \quad (2)$$

The parameter  $\delta$  governs the strength and sign of state dependence. When  $\delta = 0$ , health does not affect marginal utility.

**Frailty.** Frailty can take values between zero and one. The larger the frailty, the worse a household's health is. If a household's frailty is zero in period  $t$ , there is a positive probability that it becomes positive in the next period. I also assume that if frailty is positive in period  $t$ , it cannot go back to zero in period  $t + 1$ . This is consistent with the data: in my PSID sample, fewer than 1% of households are transitioning from positive to zero frailty. Once frailty is positive, I follow Hosseini, Kopecky, and Zhao (2022) and assume that it evolves according to the following process:

$$\log(f_t) = \kappa_t + \pi_t^f + \varepsilon_t^f, \quad (3)$$

$$\pi_t^f = \rho_f \pi_{t-1}^f + \eta_t^f, \quad (4)$$

$$\varepsilon_t^f \sim \mathbb{N}(0, \sigma_{\varepsilon^f}^2), \quad \eta_t^f \sim \mathbb{N}(0, \sigma_{\eta^f}^2), \quad \pi_0^f \sim \mathbb{N}(0, \sigma_{\pi_0^f}^2), \quad (5)$$

where  $\kappa_t$  is a deterministic component that depends on age;  $\pi_t^f$  is a persistent component, and  $\varepsilon_t^f$  is a transitory component. I assume that the shocks  $\varepsilon_t^f$  and  $\eta_t^f$  are mutually and serially uncorrelated. Following Hosseini, Kopecky, and Zhao (2022), I also assume that, when  $f_t = 0$ ,  $\pi_t^f = 0$ .

**Earnings.** Households face earnings risk during their working period. Earnings depend on age, frailty, and a persistent and transitory component as follows

$$\log y_t(f_t) = \kappa_t(f_t) + \pi_t^y + \varepsilon_t^y, \quad (6)$$

$$\pi_t^y = \rho_y \pi_{t-1}^y + \eta_t^y, \quad (7)$$

$$\varepsilon_t^y \sim \mathbb{N}(0, \sigma_{\varepsilon^y}^2), \quad \eta_t^y \sim \mathbb{N}(0, \sigma_{\eta^y}^2), \quad \pi_0^y \sim \mathbb{N}(0, \sigma_{\pi_0^y}^2), \quad (8)$$

where  $\kappa_t(f_t)$  denotes a deterministic function of age and frailty,  $\pi_t^y$  is a persistent component, and  $\varepsilon_t^y$  is a transitory component. I assume that the shocks  $\varepsilon_t^y$  and  $\eta_t^y$  are mutually and serially uncorrelated.

**Medical expenses and death.** In each period prior to death, households incur out-of-pocket medical expenses and face survival risk. I model the evolution of out-of-pocket medical expenses as follows,

$$\log m_t(f_t) = g(t, f_t) + \xi_t, \quad (9)$$

$$\xi_t \sim \mathbb{N}(0, \sigma_{\xi}^2), \quad (10)$$

where  $g(t, f_t)$  denotes a deterministic function of age and frailty and  $\xi_t$  denotes an i.i.d. shock. Medical expenses also occur for households in perfect health (i.e., zero frailty). They capture, for instance, preventive care, such as routine physicals and examinations. Households face an age-and-frailty-specific survival probability,  $s_{f,t}$ , up to the maximum age of 89.

**Government.** The government imposes taxes on income, provides Social Security benefits to retirees, and provides a means-tested transfer to needy households. I model taxes on total income as in Bénabou (2002), Heathcote, Storesletten, and Violante (2017), and Borella et al. (2023). This tax function allows for negative tax rates (and thus incorporates the Earned Income Tax Credit (EITC)) and is given by:

$$T(y_t) = y_t - (1 - \lambda)y_t^{1-\tau}, \quad (11)$$

where  $y_t$  denotes the level of total pre-tax income,  $\lambda$  captures the average level of taxation in the economy, and  $\tau$  denotes the degree of progressivity of the income tax system.

The government provides Social Security benefits after retirement. To reduce computational costs, I follow De Nardi, Fella, and Paz-Pardo (2019) and approximate Social Security benefits as a function of the last realization of earnings:

$$ss_t = ss(y_{T^{ret}-1}). \quad (12)$$

Government-provided means-tested transfers,  $b_t$ , ensure that a household's available resources are sufficient to reach a minimum consumption floor,  $\underline{c}$ :

$$b_t = \max\{0, \underline{c} + m_t(f_t) - [a_t + (1 - \lambda)[ra_t + y_t(f_t)]^{1-\tau}]\}, \quad \text{if } t < T^{ret}, \quad (13)$$

$$b_t = \max\{0, \underline{c} + m_t(f_t) - [a_t + (1 - \lambda)(ra_t + ss_t)^{1-\tau}]\}. \quad \text{if } t \geq T^{ret}. \quad (14)$$

**Timing.** Working-age households start each period with a stock of assets and draw realizations of the stochastic process for frailty, earnings, and medical expenses, and then make consumption and saving decisions. Retired households start each period with a stock of assets and Social Security benefits that remain constant until they die. They draw realizations of the stochastic processes of frailty and medical expenses and then make consumption and saving decisions.

**Recursive Formulation.** There are two value functions, one for each stage of life. The vector of state variables  $X_t$  for workers includes: age,  $t$ ; assets,  $a_t$ ; the medical expenses shock,  $\xi_t$ ; the persistent earnings component,  $\pi_t^y$ ; the transitory earnings component,  $\varepsilon_t^y$ ; the persistent frailty component,  $\pi_t^f$ ; and the transitory frailty component,  $\varepsilon_t^f$ . Workers maximize the objective function:

$$V(X_t) = \max_{c_t, a_{t+1}} \left\{ \delta(f_t) \frac{c_t^{1-\gamma}}{1-\gamma} + \beta s_{f,t} \mathbb{E}_t[V(X_{t+1})] \right\}, \quad (15)$$

subject to the intertemporal budget constraint,

$$a_{t+1} = a_t + y^n(ra_t + y_t(f_t)) - m_t(f_t) + b_t - c_t, \quad (16)$$

and Equations (3)-(5), (6)-(8), (9)-(10), (13), and a no borrowing constraint in every period,  $a_{t+1} \geq 0$ . The function  $y^n(\cdot)$  maps pre-tax total income into post-tax income using the tax function in Equation 11.

As the recursive formulation for workers makes clear, I abstract from endogenous labor supply. During working life, earnings follow the exogenous stochastic process in Equations (6)-(8), which depends on age and frailty. The model therefore captures the reduced-form effect of health on income observed in the data, while isolating how health shifts marginal utility.

The vector of state variables  $X_t$  for retirees comprises: age,  $t$ ; assets,  $a_t$ ; the medical expenses shock,  $\xi_t$ ; Social Security benefits,  $ss_t$ ; the persistent frailty component,  $\pi_t^f$ ; and the transitory frailty component,  $\varepsilon_t^f$ . Retirees maximize the objective function:

$$V(X_t) = \max_{c_t, a_{t+1}} \left\{ \delta(f_t) \frac{c_t^{1-\gamma}}{1-\gamma} + \beta s_{f,t} \mathbb{E}_t[V(X_{t+1})] \right\}, \quad (17)$$

subject to the intertemporal budget constraint,

$$a_{t+1} = a_t + ss^n(ra_t + ss_t) - m_t(f_t) + b_t - c_t, \quad (18)$$

and Equations (3)-(5), (9)-(10), (12), (14), and a no borrowing constraint in every period,  $a_{t+1} \geq 0$ . The function  $ss^n(\cdot)$  maps pre-tax total income into post-tax income using the tax function in Equation 11.

## IV Data

**Dataset.** The PSID is a longitudinal survey representative of the U.S. population, conducted annually since 1968 and biennially since 1997. Starting in 2003, the PSID collected detailed health information and, starting in 2005, it covered almost all of the consumption

categories considered in the Consumer Expenditure Survey (CEX). I therefore use each biennial wave of the PSID between 2005 and 2019 because, during this sample period, the dataset contains detailed information on health and medical conditions, labor and non-asset income, wealth, and consumption. To be consistent with my model, I consider households whose heads are between 25 and 89 years old. Appendix A provides details about my data and sample selection.

I measure non-medical consumption as the sum of household expenses on food at and away from home, utilities, phone bills, internet bills, transportation (excluding car loans, lease payments, and down payments), trips and vacations, entertainment and recreation, donations to charity, and clothing. I convert nominal earnings, medical expenses, and consumption into real quantities using the Consumer Price Index for Urban Consumers (CPI-U) and 2018 as my base year.

**Measuring health.** I measure health using the frailty index. The frailty index captures the idea that, as people age, they become increasingly burdened by adverse health events (such as chronic diseases or temporary ailments), which I refer to as deficits. As such, the frailty index is an objective measure of bad health. It has been used extensively in the medical and gerontology literature, and it has been shown to be an excellent predictor of health, mortality, medical expenses, and the probability of becoming a disability insurance recipient (See, among others, Hosseini, Kopecky, and Zhao (2022), Nygaard (2021), Russo et al. (2024), and Borella et al. (2025)).

To construct the frailty index for the households in my sample, I follow Searle et al. (2008) and include deficits belonging to the following categories: difficulties with activities of daily living (ADL) and instrumental ADL (IADL) (such as difficulty dressing, bathing, and walking); diagnosed diseases (such as diabetes, cancer, and arthritis); cognitive impairments and mental health measures (such as memory loss and psychological problems); and lifestyle habits (such as obesity, smoking, and excessive alcohol consumption).<sup>3</sup> In total, I consider

---

<sup>3</sup>I follow the medical literature in classifying lifestyle habits as health deficits. The American Medical Association recognized obesity as a chronic disease in 2013 and classified alcohol addiction as an illness as early as 1956. Medical studies such as Bernstein and Toll (2019) also treat smoking as a chronic condition. Excluding these deficits from frailty would therefore be difficult to justify on medical grounds. Nevertheless, in Appendix Appendix B, I construct an alternative frailty index that excludes the four deficits related to

29 possible deficits to construct the frailty index. Each can take a value of either zero or one, depending on whether the individual currently has a specific deficit or not. Table A-2 in Appendix B reports the complete list of deficits.

I aggregate deficits into the frailty index by taking the (equally weighted) sum of a person’s deficits at a given age and dividing it by the total number of deficits I consider. Hosseini, Kopecky, and Zhao (2022) and Russo et al. (2024) show that applying other weighting schemes (such as Principal-Component-Analysis-derived weights) to the deficits composing the frailty index does not change the dynamics of health evolution nor improve the predictive power of frailty for key health and economic outcomes. I then construct a household’s frailty as the average frailty index of each member.

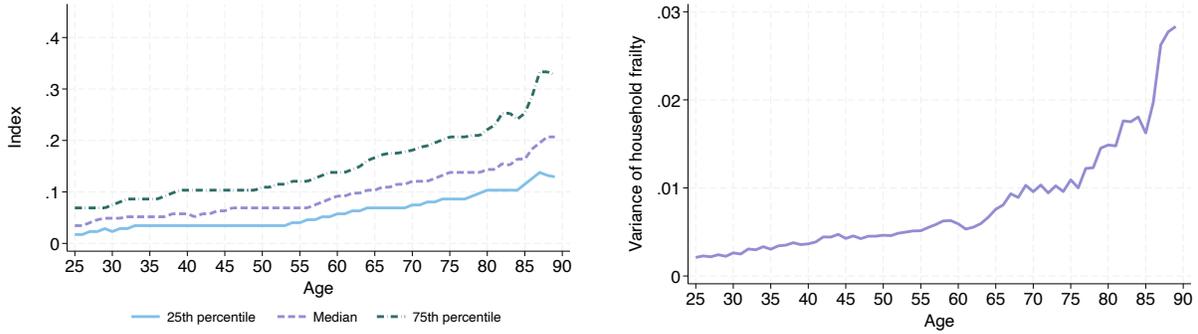
Household frailty has a mean of 0.09 and a median of 0.07. Panel (a) of Figure 1 reports the distribution of household frailty by age and shows that the median, 25th, and 75th percentiles increase with age. Panel (b) reports the variance of frailty by age and shows that this increases with age and is particularly high after age 75. Panel (c) of Figure 1 displays the fraction of households with positive frailty by age and shows that health risk is present at all ages. In particular, it indicates that over 80% of 25-year-old households have positive frailty and that this share increases rapidly with age. To understand why frailty is so prevalent, even at younger ages, Table 1 displays the prevalence of each health deficit for household heads between 25 and 29. In Appendix Appendix B, I report the prevalence of the same deficits for household heads of ages 45 to 49, 65 to 69, and 85 to 89, respectively. These tables show that smoking and obesity are the most common deficits for younger people, while high blood pressure and arthritis are the most common for older people.<sup>4</sup>

---

obesity, smoking, and excessive alcohol consumption. Doing so has a modest effect on frailty dynamics and does not significantly alter the estimated consumption response to a transitory health shock.

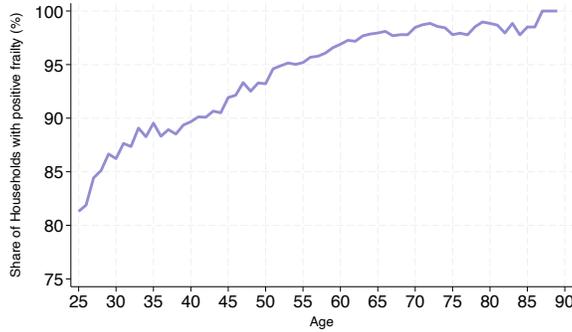
<sup>4</sup>These numbers are consistent with statistics from the Centers for Disease Control and Prevention (CDC). In particular, the CDC documents that, in 2020, 14.1% of people aged 25 to 44 smoked cigarettes (CDC (2022b)), while about 40% of people aged 20 to 39 were obese (CDC (2022a)). Moreover, the CDC reports that, in 2018, 74.5% of people over 60 had high blood pressure (CDC (2020)), while, between 2013 and 2015, about 50% of people over 65 had arthritis (CDC (2021)).

Figure 1: Frailty statistics by age



(a) Percentiles

(b) Variance



(c) Share of households with positive frailty

Notes: Frailty statistics by age: 25th, 50th, 75th percentiles (a); variance (b); and share of households with positive frailty (c). Each statistic is smoothed using a 3-year moving average. PSID, 2005-2019.

Table 1: Prevalence of health deficits

Deficit	Prevalence	Deficit	Prevalence
BMI $\geq 30$	0.289	Diagnosed with heart disease	0.013
Smokes now	0.256	Diff. getting in/out of bed	0.011
Has ever smoked cigarettes	0.175	Diagnosed with loss of memory or mental ability	0.012
Diagnosed with asthma	0.168	Diff. light housework	0.010
Diagnosed with other serious chronic conditions	0.153	Diagnosed with stroke	0.009
Diagnosed with psychological problems	0.128	Diff. shopping for personal items	0.007
Diagnosed with high blood pressure	0.124	Diff. dressing	0.006
Excessive alcohol drinking	0.059	Diff. preparing meals	0.006
Diagnosed with lung disease	0.056	Diff. bathing/showering	0.005
Diagnosed with arthritis	0.045	Diff. getting outside	0.005
Diagnosed with diabetes	0.022	Diff. eating	0.004
Diff. heavy housework	0.021	Diff. using the toilet	0.003
Diff. walking	0.018	Diagnosed with heart attack	0.003
Diagnosed with cancer	0.016	Diff. using the phone	0.001
Diff. managing money	0.014		

Notes: Household heads between 25 and 29. PSID, waves 2005-2019.

## V Calibration

I use a two-step strategy similar to that of Gourinchas and Parker (2002) and De Nardi, French, and Jones (2010). In the first step, I estimate the parameters that I can cleanly identify outside my model. For example, I estimate the frailty process from the PSID and fix the discount factor and risk aversion to values commonly used in the literature.

In the second step, I calibrate the consumption floor,  $\underline{c}$ , and the effect of health on preferences,  $\delta$ , to match the consumption response to transitory earnings and health shocks.

### V.1 First-step calibration and estimation

**Frailty process.** I model the probability that the household's frailty is zero at each age using a probit model as in Hosseini, Kopecky, and Zhao (2022):

$$\text{Prob}(f_t = 0|X_t) = \Phi(X_t'\alpha), \quad (19)$$

where  $\Phi$  is the c.d.f. of a standard normal distribution and  $X_t$  is a set of regressors. Here,  $X_t$  contains family size, education level, cohort effects, and a second-order polynomial in household age. Table A-6 in Appendix C reports the probit regression results.

The probability that a household has zero frailty, conditional on having zero frailty in the previous period, is:

$$\text{Prob}(f_t = 0|f_{t-1} = 0) = \frac{\text{Prob}(f_t = 0|X_t)}{\text{Prob}(f_{t-1} = 0|X_{t-1})} = \frac{\Phi(X_t'\alpha)}{\Phi(X_{t-1}'\alpha)}. \quad (20)$$

Thus, the probability that a household has zero frailty is given by Equation (19) at age 25; given by Equation (20) at ages older than 25 if frailty is zero in the previous period; and zero otherwise. Figure A-2 in Appendix C displays the share of households with zero frailty in the raw data and those predicted by my probit regression.

For households with positive frailty, I estimate the deterministic component ( $\kappa_t$  in Equation (3)) by regressing log-frailty on family size, education level, cohort effects, and a second-order polynomial in age. Using the residuals from this regression, I estimate the autoregressive coefficient,  $\rho_f$ , the variance of the transitory shock,  $\sigma_{\varepsilon_f}^2$ , the variance of the shock to

the persistent component,  $\sigma_{\eta f}^2$ , and the variance of the initial persistent component,  $\sigma_{\pi_0^f}^2$ . I identify them using the variances and covariances of the residuals and estimate them using GMM, targeting the variance and autocovariance age profiles in the PSID (see Appendix C for identification restrictions and estimation details). Table A-7 in Appendix C reports the results.

**Survival probabilities.** I estimate age- and frailty-specific two-year survival probabilities for household heads. Specifically, I run a logistic regression of a binary indicator of survival between waves  $t$  and  $t + 1$  on lagged frailty, education, family size, cohort effects, and a second-order polynomial in age. Table A-8 in Appendix C reports the estimation results. I then compute the average survival probabilities by age and confirm the finding in French (2005) that the PSID overestimates survival probabilities. Hence, I adjust them so that my estimated average survival probabilities match those reported by the Social Security Administration in the 2019 life tables.

**Earnings process.** Earnings include labor earnings, the labor part of business income, and farm income. When married, household earnings are the sum of each spouse’s earnings. I estimate the earnings process for households aged 25 to 61 with positive labor earnings. I estimate the deterministic function  $\kappa_t(f)$  in Equation (6) by regressing the logarithm of earnings on frailty, family size, education level, cohort effects, and a second-order polynomial in age. The left panel Table A-9 in Appendix C displays the estimation results. Using the residuals from this regression, I estimate the autoregression coefficient  $\rho_y$ , the variance of the transitory shock,  $\sigma_{\varepsilon^y}^2$ , the variance of the shock to the persistent component,  $\sigma_{\eta^y}^2$ , and the variance of the initial persistent component,  $\sigma_{\pi_0^y}^2$ , using GMM, targeting the variance and autocovariance age profiles in the PSID. Appendix C presents details on the identification and estimation. The right panel of Table A-9 in Appendix C shows the estimated variances of the earnings shocks.

**Out-of-pocket medical expenses.** Medical expenses are the sum of what households spend out-of-pocket for hospital and nursing home stays, doctor visits, prescription drugs, and insurance premiums. I replace values of medical expenses equal to zero with \$100. I estimate

the deterministic function  $g(t, f_t)$  in Equation (9) by regressing the logarithm of medical expenses on frailty, family size, education level, cohort effects, and a second-order polynomial in household age. Column (1) in Table A-10 in Appendix C reports the estimation results for this regression. To estimate the variance of the i.i.d. shock,  $\sigma_\xi^2$ , I regress the squared residuals from the regression above on the same covariates. Column (2) of Table A-10 in Appendix C reports these estimation results. I then compute the predicted values from this regression and their variance, which provides the estimate for the variance of the i.i.d. shock.

**Fixed parameters.** Table 2 summarizes my first-step parameters, including those that I set to commonly adopted values in the literature. I use the tax function parameters estimated by Borella et al. (2023) for 2017 (their last available data point). I set the interest rate to two percent following Paz-Pardo (2022). I set the coefficient of relative risk aversion and the discount factors to standard values of 1.5 (as in Voena (2015)) and 0.99, respectively. I square the annual value to obtain the biennial discount factor, as in Kydland and Prescott (1982).

Table 2: First-Step Parameters

Parameter	Value	Description	Source
<i>Preference Parameters</i>			
$\beta$	0.99	Annual Discount factor	standard
$\gamma$	1.5	Risk Aversion	Voena (2015)
<i>Frailty Process</i>			
$\kappa_t$	see text	Deterministic component	PSID
$\rho_f$	0.89	Autoregressive coefficient	PSID
$\sigma_{\varepsilon^f}^2$	0.01	Variance of transitory shock	PSID
$\sigma_{\eta^f}^2$	0.09	Variance of persistent shock	PSID
$\sigma_{\pi_0^f}^2$	0.34	Variance of initial persistent component	PSID
<i>Earnings Process</i>			
$\kappa_t(f_t)$	see text	Deterministic function of age and frailty	PSID
$\rho_y$	0.99	Autoregressive coefficient	PSID
$\sigma_{\varepsilon^y}^2$	0.09	Variance of transitory shock	PSID
$\sigma_{\eta^y}^2$	0.01	Variance of persistent shock	PSID
$\sigma_{\pi_0^y}^2$	0.36	Variance of initial persistent component	PSID
<i>Medical Expenses and Death</i>			
$g(f_t, t)$	see text	Deterministic function of age and frailty	PSID
$\sigma_\xi^2$	0.04	Variance of shock to medical expenses	PSID
$s_{f,t}$	see text	Age-and-frailty-specific survival probabilities	PSID
<i>Government and Interest Rate</i>			
$\lambda$	2	Average level of income taxation	Borella et al. (2023)
$\tau$	-0.07	Progressivity of the income tax system	Borella et al. (2023)
$r$	0.02	Interest rate	Paz-Pardo (2022)

## V.2 Second-step calibration

I calibrate the consumption floor and the effect of bad health on the marginal utility of consumption to match the degree of self-insurance against transitory earnings and frailty shocks in the data.

**Measuring self-insurance.** I follow a long-standing tradition in the consumption insurance literature and measure self-insurance with pass-through coefficients.<sup>5</sup> As described in the seminal contribution of Blundell, Pistaferri, and Preston (2008), the pass-through coefficient of an idiosyncratic shock  $x_t$  is the ratio of the covariance between log-consumption growth and the shock and the variance of the shock:

$$\phi^x = \frac{\text{cov}(\Delta \log c_t, x_t)}{\text{var}(x_t)}. \quad (21)$$

Pass-through coefficients capture the share of the variance of a shock that translates into consumption growth. If households had access to full insurance against idiosyncratic shocks (as would be the case if markets were complete), the pass-through coefficients would be zero. If there were no insurance, the pass-through coefficients would be one, as consumption would react one-to-one to idiosyncratic shocks. As shown by Blundell, Pistaferri, and Preston (2008) and Blundell et al. (2024), households in the United States have only partial insurance against income and health shocks, resulting in pass-through coefficients between zero and one.

Estimating pass-through coefficients from the data is challenging because shocks are not observable. For example, the PSID records information on earnings and frailty but not on earnings and frailty shocks. Therefore, after “detrending” them from observable characteristics, I use moments on observable consumption, earnings, and frailty to estimate the pass-through coefficients. Appendix D presents more details about the “detrending” procedure.

---

<sup>5</sup>A few notable papers using pass-through coefficients to measure consumption insurance are Blundell, Pistaferri, and Preston (2008), Kaplan and Violante (2010), Blundell, Pistaferri, and Saporta-Eksten (2016), Blundell, Pistaferri, and Saporta-Eksten (2018), Wu and Krueger (2018), and Blundell et al. (2024)

I apply a similar strategy to that used by Kaplan and Violante (2010) to identify the pass-through coefficients of transitory earnings and frailty shocks (Appendix G describes the details for the pass-through coefficients of persistent shocks.) First, I define the quasi difference of log earnings as  $\tilde{\Delta} \log y_t = \log y_t - \rho_y \log y_{t-1}$  and the quasi difference of log frailty as  $\tilde{\Delta} \log f_t = \log f_t - \rho_f \log f_{t-1}$ . Notice that I have estimated  $\rho_y$  and  $\rho_f$  in Section V.1. Second, using the processes for earnings and frailty and assuming that shocks are mutually uncorrelated, one can show that

$$\begin{aligned} \text{cov}(\Delta \log c_t, \tilde{\Delta} \log y_{t+1}) &= -\rho_y \text{cov}(\Delta \log c_t, \varepsilon_t^y), \\ \text{cov}(\tilde{\Delta} \log y_t, \tilde{\Delta} \log y_{t+1}) &= -\rho_y \text{var}(\varepsilon_t^y), \\ \text{cov}(\Delta \log c_t, \tilde{\Delta} \log f_{t+1}) &= -\rho_f \text{cov}(\Delta \log c_t, \varepsilon_t^f), \\ \text{cov}(\tilde{\Delta} \log f_t, \tilde{\Delta} \log f_{t+1}) &= -\rho_f \text{var}(\varepsilon_t^f). \end{aligned}$$

Therefore, the pass-through coefficients to transitory shocks are identified as:

$$\phi_\varepsilon^y = \frac{\text{cov}(\Delta \log c_t, \varepsilon_t^y)}{\text{var}(\varepsilon_t^y)} = \frac{\text{cov}(\Delta \log c_t, \tilde{\Delta} \log y_{t+1})}{\text{cov}(\tilde{\Delta} \log y_t, \tilde{\Delta} \log y_{t+1})}, \quad (22)$$

$$\phi_\varepsilon^f = \frac{\text{cov}(\Delta \log c_t, \varepsilon_t^f)}{\text{var}(\varepsilon_t^f)} = \frac{\text{cov}(\Delta \log c_t, \tilde{\Delta} \log f_{t+1})}{\text{cov}(\tilde{\Delta} \log f_t, \tilde{\Delta} \log f_{t+1})}. \quad (23)$$

I estimate the pass-through coefficients of transitory shocks using Equation (22) and Equation (23). Appendix D provides details on the estimation procedure.

**Calibration procedure.** I calibrate the consumption floor,  $\underline{c}$ , and the effect of bad health on marginal utility,  $\delta$ . To solve and simulate the model, I follow Gourinchas and Parker (2002) and French (2005) and fix the cohort to the middle one in the data, fix family size to the average family size, and fix the education level to high school graduate. I keep these values constant over the life cycle, as in French (2005). Then, the calibration procedure proceeds as follows. First, given an initial guess for the two parameters to be calibrated, I solve the life-cycle model and obtain optimal decision rules for consumption and savings.

Second, I use the optimal decision rules to simulate the life-cycle choices of households. Third, following Commault (2022), I detrend log-earnings, log-frailty, and log-consumption by regressing them on age dummies. Fourth, using these detrended data, I compute the pass-through coefficients for transitory earnings and health shocks. Fifth, I compute the objective function for my calibration as the squared difference between the pass-through coefficients in the model and the data. Finally, using the Nelder-Mead algorithm, I search for the combination of  $\delta$  and  $\underline{c}$  that minimizes my objective function. Appendix F provides more details on the model solution and simulation.

**Identification.** In a non-linear model like mine, all parameters potentially affect all moments. In the Introduction, I provide some intuition on what moments in the data help identify my parameters of interest. Appendix E provides more details and shows the identification argument for the effect of health on preferences using the structure of the model.

**Calibration results.** The second column of Table 3 reports the pass-through coefficients to transitory earnings and frailty shocks that I estimate from the PSID. The third column shows the corresponding coefficients in my simulated data, and the fifth column shows the values of my calibrated parameters.

Table 3: Targeted moments, model fit, and parameter values

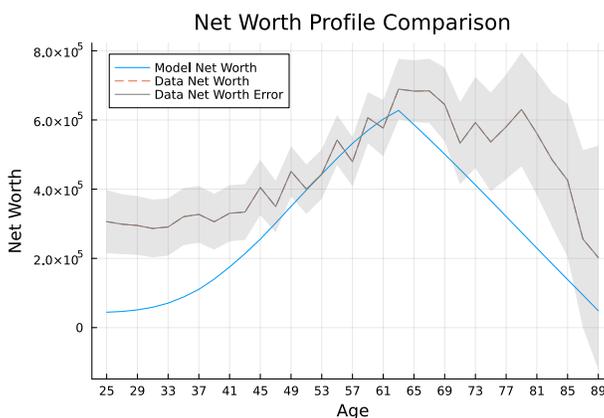
Moment	Data	Model	Parameter	Value
$\phi_{\varepsilon}^y$	0.125	0.122	Consumption floor, $\underline{c}$	\$2,566
$\phi_{\varepsilon}^f$	-0.103	-0.103	Effect of bad health on marginal utility, $\delta$	-0.89

The first row of Table 3 shows that the pass-through coefficient of a transitory earnings shock,  $\phi_{\varepsilon}^y$ , is 0.125. This means that a 10% increase in earnings due to a transitory shock leads to a rise in consumption of 1.25%. This finding is in line with the results of Blundell, Pistaferri, and Preston (2008) and Blundell et al. (2024). My model matches the pass-through of a transitory earning shock and predicts an annual consumption floor of \$2,566. The second row of Table 3 shows that the pass-through coefficient of a transitory frailty shock,  $\phi_{\varepsilon}^f$ , is -0.103. This means that a 10% increase in frailty due to a transitory shock

leads to a 1% decrease in consumption.<sup>6</sup> My model replicates the pass-through of transitory frailty shocks and implies a calibrated value of  $\delta$  of -0.89. The fact that  $\delta$  is negative means that the marginal utility of consumption decreases as health worsens. This finding is in line with the results of, among others, De Nardi, French, and Jones (2010), Finkelstein, Luttmer, and Notowidigdo (2013), Koijen, Van Nieuwerburgh, and Yogo (2016), and Blundell et al. (2024).

**Untargeted Moments.** To validate my model, I compare the evolution of net worth in the model to that in the data. In the PSID, I measure a household’s net worth as the sum of all assets minus all liabilities. In particular, I define it as the sum of the equity in farms and businesses; transaction accounts (such as savings accounts, money market funds, certificates of deposits, government bonds, and treasury bills); equity in real estate, stock, vehicles, and IRAs; the value of home equity (calculated as home value minus remaining mortgage); net of total debt. Figure 2 shows the comparison between the model-implied net worth and the one from the PSID. Despite starting from a lower initial value, the model captures the accumulation of assets until retirement and does particularly well between the ages of 45 and 65. The model also captures the decumulation of assets after retirement, though it predicts a faster pace, consistent with the absence of a bequest motive.

Figure 2: Net worth: model vs. data



<sup>6</sup>A potential concern is that transitory health shocks may affect consumption through labor-force exit. Restricting attention to households whose head reports working in periods  $t - 1$ ,  $t$ , and  $t + 1$  yields a pass-through estimate of  $-0.064$  (s.e. 0.055), compared to  $-0.103$  in the full sample. The persistence of a negative response when employment status is held fixed indicates that labor-force transitions are not the primary driver of the consumption response used for identification. Table A-12 in Appendix D reports details.

Next, I compare the pass-through coefficients against persistent earnings and frailty shocks in the model and the data. As I argued in the Introduction, identifying the pass-through coefficients requires stringent assumptions on consumption, which are violated in the observed and the model-generated data. However, I compare the estimates in the data and the ones from the model to evaluate my model’s performance, even though they are biased measures of the true degree of self-insurance against persistent shocks. Table 4 reports the results of this comparison and shows that my model generates results that are qualitatively similar to those in the data.<sup>7</sup> In particular, a persistent positive earnings shock generates a consumption increase both in the data and in the model. In terms of size, Table 4 shows that my model generates a pass-through of persistent earnings shocks that is larger than that in the data (but very close to the estimates of Kaplan and Violante (2010)). Conversely, it implies a pass-through of persistent health shocks that is lower than what is observed (albeit without statistical significance) in the data. The bias in the model-generated pass-through coefficient to persistent shocks could be due to the identification strategy issues, leading to biased estimates both in the model and in the data. Moreover, the bias in the pass-through to a persistent frailty shock could also be due to misspecification of the long-term effects of bad health, such as its effects on survival.

Table 4: Pass-through coefficients of persistent shocks

$\phi_{\eta}^y$		$\phi_{\eta}^f$	
Data	Model	Data	Model
0.37***	0.76***	0.01	-0.10***
(0.05)	(0.04)	(0.03)	(0.02)
7,414	79,957	8,441	150,000

<sup>7</sup>Appendix G provides more details on estimating the pass-through coefficients of persistent shocks and the assumptions needed for identification.

## VI Health-Dependent Preferences, Consumption, and Savings

In this section, I use my calibrated model to assess the quantitative effects of health-dependent preferences on consumption and savings. This channel has received little attention in the quantitative life-cycle literature. For example, Scholz, Seshadri, and Khitatrakun (2006) studies retirement saving adequacy while incorporating medical expense risk, but abstracts from health-driven changes in marginal utility.

To quantify the effects of health-dependent preferences on consumption and savings, I solve the model using the baseline calibration and setting  $\delta = 0$ , thereby removing the relationship between health and marginal utility. I then simulate the life cycles of 50,000 households and compare their consumption and savings with those from my baseline calibration.

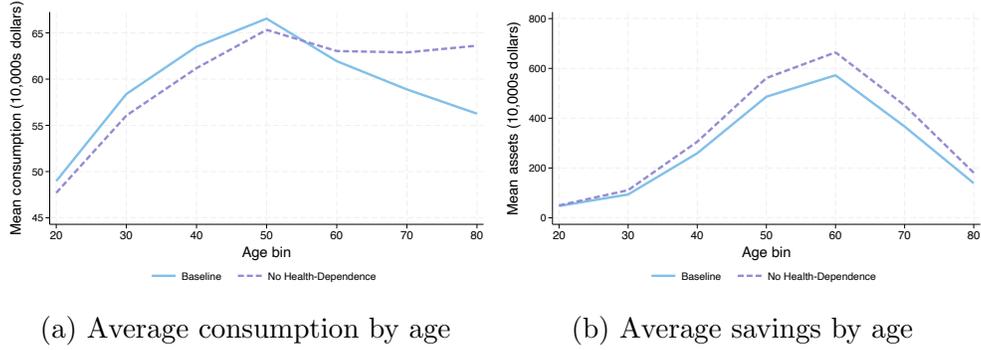
Health-dependent preferences affect optimal consumption and savings over the life cycle. Figure 3 plots the life-cycle profile of average consumption and savings by 10-year age bins. Panel (a) shows that average consumption is lower without health-dependent preferences before 50 but higher at older ages. For example, without health-dependent preferences, households consume about 2.6% less in their twenties but 13% more in their eighties. Panel (b) shows that average savings are higher at every age without health-dependent preferences. In particular, the increase in savings ranges from 4.7% when households are in their twenties to 30% in their eighties. These results show that health-dependent preferences should be carefully considered when studying the patterns and determinants of savings.

The observed consumption and savings patterns are consistent with the deterioration of health over the life cycle. As health worsens with age, shutting down state dependence raises marginal utility in old age relative to the baseline. As a result, optimal consumption increases in old age relative to the baseline. To sustain higher consumption at older ages, households must save more and give up consumption when young.<sup>8</sup>

---

<sup>8</sup>Using a stylized two-period model, Finkelstein, Luttmer, and Notowidigdo (2013) shows that an adverse effect of bad health on the marginal utility of non-medical consumption reduces the optimal level of savings for retirement, which is consistent with what I find.

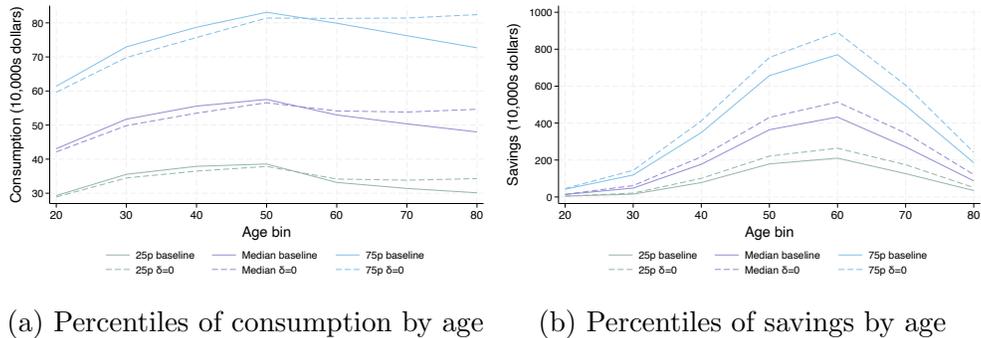
Figure 3: Consumption and savings profiles



Notes: Panel (a) shows average consumption by 10-year age bin for the baseline calibration and the counterfactual experiment in which  $\delta = 0$ . Panel (b) shows the analogous figure for savings.

The patterns described above apply across the consumption and savings distributions. Figure 4 displays the 25th, 50th, and 75th percentiles of consumption and savings by 10-year age bins with and without health-dependent preferences. Panel (a) shows that households consume more without health-dependent preferences at older ages and less at younger ones. The magnitude of the effect is similar across percentiles; for example, in their eighties, consumption is roughly 14 percent higher without state dependence. Panel (b) shows the 25th, 50th, and 75th percentiles of savings by age. This figure shows that all households save more without health-dependent preferences at all ages. For example, in their thirties, savings rise by roughly 20 to 35 percent across percentiles.

Figure 4: Percentiles of consumption and savings



Notes: Panel (a) shows the 25th, 50th, and 75th percentile of consumption by 10-year age bin for the baseline calibration and the counterfactual experiment in which  $\delta = 0$ . Panel (b) shows the analogous figure for savings.

## VII Health-Dependent Preferences and Self-Insurance

I next quantify how health-dependent preferences affect households' ability to self-insure against health shocks. To do so, I compute the pass-through coefficients of transitory health shocks in my baseline model and compare them with those from a counterfactual in which health does not affect the marginal utility of consumption (and  $\delta = 0$ ). With the exception of Blundell et al. (2024), this channel has not been studied in the context of quantitative life-cycle models.

Table 5 reports the pass-through coefficient of transitory frailty shocks in my PSID sample, baseline model, and a model without health-dependent preferences. This table shows that the model without health-dependent preferences predicts that a 10% increase in frailty generated by a transitory shock results in a 0.6% decrease in consumption. This response is 38% smaller than both the empirical estimate and the baseline model prediction. The counterfactual, therefore, implies a smaller consumption response (and thus a higher degree of self-insurance) against adverse health shocks. This is consistent with the higher saving rates in the counterfactual economy, but inconsistent with the consumption responses observed in the PSID. Shutting down state dependence, therefore, mechanically attenuates the consumption response to health shocks, leading the model to overstate households' ability to self-insure.

Table 5: Comparison of pass-through coefficients

Moment	Data	Baseline $\delta = -0.89$	No health-dependent preferences $\delta = 0$
$\phi_{\varepsilon}^f$	-0.103	-0.103	-0.064

Notes: Pass-through coefficients of transitory frailty shocks in the baseline calibration,  $\delta = -0.89$ , and counterfactual experiment in which  $\delta = 0$ .

## VIII Health-Dependent Preferences and Government Insurance

In this section, I analyze the welfare effects of reforming means-tested government insurance (MTGI) with and without health-dependent preferences. Accurately measuring households' valuation of these programs is crucial to evaluating reforms. A related public finance literature studies the value of disability insurance (Deshpande and Lockwood, 2022). Because MTGI includes programs targeted at the poor and the unhealthy (such as Medicaid and Supplemental Security Income), health-dependent preferences may play a key role in determining its value.

I compute the welfare changes associated with MTGI reforms using the compensating variation. In particular, I follow De Nardi, French, and Jones (2016) and McGee (2021) and define the compensating variation as the immediate payment after the reform that would make households as well off as before the reform. I compute the compensating variation at age 25 (the initial age in my model and simulations) and define it as  $\chi_{25}(a_{25}, \xi_{25}, \pi_{25}^y, \varepsilon_{25}^y, \pi_{25}^h, \varepsilon_{25}^h)$  solving:

$$V_{25}(a_{25}, \xi_{25}, \pi_{25}^y, \varepsilon_{25}^y, \pi_{25}^h, \varepsilon_{25}^h | \text{Baseline}) = V_{25}(a_{25} + \chi_{25}, \xi_{25}, \pi_{25}^y, \varepsilon_{25}^y, \pi_{25}^h, \varepsilon_{25}^h | \text{Reform}),$$

where  $V_{25}(\cdot)$  is the age 25 value function for a given set of state variables. As argued in McGee (2021), the compensating variation is an ex-ante measure that incorporates the mechanical and behavioral responses to a reform.

Health-dependent preferences affect the household valuation of MTGI. Table 6 shows the compensating variation associated with a reform that reduces the consumption floor by 50%. First, households place a higher value on government insurance without health-dependent preferences. In particular, the first row of Table 6 shows that for all 25-year-olds (who earn \$51,116 on average and have a frailty index of 0.06), the compensating variation associated with a 50% reduction in the consumption floor is about 6% higher without health-dependent preferences.

Table 6: Household valuation of MTGI reform

Group	Average Earnings	Average Frailty	50% reduction in $c$	
			Baseline $\delta = -0.89$	No HDP $\delta = 0$
All 25-year olds	51,116	0.06	89	94
Bottom 5% earnings	10,048	0.09	609	658
Top 5% earnings	167,473	0.04	1	1
Bottom 5% frailty	63,467	0.00	82	86
Top 5% frailty	32,528	0.19	102	107

Notes: Columns 4 and 5 report the compensating variation associated with a 50% reduction in the consumption floor. I report the compensating variation in the baseline calibration,  $\delta = -0.89$ , and without health-dependent preferences,  $\delta = 0$ .

Second, government insurance is more valuable for low earners. In the second and third rows of Table 6, I compare households in the bottom and top 5% of the earnings distribution and show that the welfare effects are larger for low earners, who are, on average, also more unhealthy. In turn, households in the top 5% of the earnings distribution place no value on the consumption floor, which is consistent with the fact that their resources are high enough for the floor to never bind. For low-earners, the household valuation of government insurance is larger without health-dependent preferences.

Third, government insurance is more valuable for sicker households. In the fourth and fifth rows of Table 6, I compare relatively healthy households (i.e., in the bottom 5% of the frailty distribution) with rather unhealthy ones (i.e., in the top 5% of the frailty distribution). My results show that sicker households (who, on average, also earn less than healthier ones) value government insurance more than healthier households do. In particular, the compensating variation is about 24% larger for households in the top 5% of the frailty distribution. The value of government insurance is higher in all cases without health-dependent preferences.

## IX Conclusions

This paper studies how health affects the marginal utility of consumption and the implications for saving, self-insurance, and the valuation of government insurance. I develop a

life-cycle model in which health affects survival, earnings, medical expenses, and the marginal utility of consumption, and I calibrate it using the PSID for quantitative analysis.

I find that bad health reduces the marginal utility of consumption and that health-dependent preferences reduce old-age consumption and lower savings throughout the whole life cycle. I also show that a model without health-dependent preferences overstates the degree of self-insurance against transitory health shocks: it predicts a smaller consumption response to bad health shocks compared to the data and the baseline model.

Health-dependent preferences also affect the household valuation of means-tested government insurance programs such as Medicaid and Supplemental Security Income. In particular, I show that households value these programs more when bad health does not affect preferences. These findings imply that the design of taxes and transfers should account for state dependence in preferences. Ignoring this channel may lead to biased assessments of both household insurance and the value of social insurance.

## APPENDIX

### A PSID data and sample selection

**The Panel Study of Income Dynamics.** The Panel Study of Income Dynamics (PSID) is a longitudinal survey of US families conducted by the University of Michigan. It was an annual survey between its inception in 1968 and 1997, and has been biennial since then. The original 1968 PSID sample contained a nationally representative sample of 2,930 households and a sample of 1,872 low-income families (the SEO subsample). The PSID follows the original 1968 families and any family member who moves out of them.

The PSID has been recording rich information on family income and wealth dynamics since 1968. Throughout the years, it has added information on respondents' social, demographic, economic, and health characteristics. In particular, until 1997, it collected only information on food consumption. Starting in 1999, it expanded its consumption measures, and since 2005, it has covered almost all the consumption categories measured by the Consumption Expenditure Survey (CEX). Moreover, in 2003, the PSID expanded its health-related questions and started recording information on specific medical conditions, ADLs, and IADLs. Johnson et al. (2018) provides a detailed description of the PSID and its changes over the last fifty years.

**Sample selection.** Table A-1 describes my sample selection. I use every biennial wave of the PSID between 2005 and 2019 and obtain an initial sample of 247,871 individual-wave observations. First, I focus on household heads. The PSID records health variables only for household heads and their spouse. Thus, I have to exclude all family members other than the two spouses from my sample. Then, household heads respond to questions about their own and their spouse's health and labor earnings, as well as total household consumption, medical expenses, and wealth. Thus, in my sample, I only keep household heads and link information on their spouse when one is present.

Then, I restrict my attention to the core sample of the PSID.<sup>9</sup> This leaves me with 42,788 observations. I remove households that appear only once in the survey. The resulting sample consists of 41,259 observations. To be consistent with my model, I focus on households whose head is between 25 and 89 years old. Then, I drop observations with missing information on frailty, labor earnings, medical expenses, wealth, family size, and the head’s education. The resulting sample contains 33,992 observations. After removing observations with missing information, I remove outliers. To do so, I first drop observations with consumption or labor earnings smaller than 50 dollars (in 2018 terms). Finally, to remove further outliers, I drop observations in the first and 99th percentiles of the change in log consumption, log income, log medical expenses, and level of wealth. The final sample consists of 28,560 observations.

Sample Selection	Selected out	Selected in
Waves 2005 - 2019		247,871
Heads only	176,696	71,175
PSID core sample	28,387	42,788
Interview in subsequent year	1,529	41,259
Age between 25 and 89	2,580	38,679
Missing key variables	4,687	33,992
Remove outliers	1,954	32,038
Drop top and bottom outliers	3,478	28,560

Table A-1: Sample Selection, PSID waves 2005 - 2019.

## B The frailty index

Table A-2 presents the complete list of deficits I use to construct the frailty index in my sample. I use 29 deficits in total. Compared to Hosseini, Kopecky, and Zhao (2022), I add alcohol consumption as a deficit. I follow the definition of the National Institute on Alcohol Abuse and Alcoholism and assign a value of one to the excessive drinking deficit if the respondent drinks every day or several times a week and, when they drink, they have more than four drinks for a man and more than three drinks for a woman. Table A-3 reports

<sup>9</sup>As discussed in Haider (2001) and Paz-Pardo (2022), the SRC subsample is a random sample, and therefore, sample weights are not needed. This is standard practice in the literature. See, for example, Blundell, Pistaferri, and Preston (2008), Heathcote, Storesletten, and Violante (2014), Blundell, Pistaferri, and Saporta-Eksten (2016), and Arellano, Blundell, and Bonhomme (2017).

the prevalence of these deficits for household heads between the ages of 45 and 49, 65 and 69, and 85 and 89.

Table A-2: Deficits used to construct the frailty index

Variable	Value	Variable	Value
<i>Some difficulty with ADL/IADLs:</i>			
Eating	Yes=1, No=0	Diabetes	Yes=1, No=0
Dressing	Yes=1, No=0	Cancer	Yes=1, No=0
Getting in/out of bed or chair	Yes=1, No=0	Lung disease	Yes=1, No=0
Using the toilet	Yes=1, No=0	Heart disease	Yes=1, No=0
Bathing/Showering	Yes=1, No=0	Heart attack	Yes=1, No=0
Walking	Yes=1, No=0	Stroke	Yes=1, No=0
Using the telephone	Yes=1, No=0	Arthritis	Yes=1, No=0
Managing money	Yes=1, No=0	Asthma	Yes=1, No=0
Shopping for personal items	Yes=1, No=0	Loss of memory or mental ability	Yes=1, No=0
Preparing meals	Yes=1, No=0	Psychological problems	Yes=1, No=0
Heavy housework	Yes=1, No=0	Other serious chronic conditions	Yes=1, No=0
Light housework	Yes=1, No=0	<i>Other conditions</i>	
Getting outside	Yes=1, No=0	BMI $\geq$ 30	Yes=1, No=0
<i>Ever had one of the following conditions:</i>			
High blood pressure	Yes=1, No=0	Has ever smoked	Yes=1, No=0
		Smokes now	Yes=1, No=0
		Excessive alcohol drinking	Yes=1, No=0

Notes: For the “Ever had one of the following conditions” variables, I make the following adjustment: If an individual reports one of these conditions, I assign a value of 1 to that deficit in every wave after the first report.

Table A-3: Prevalence of health deficits

	Prevalence		Prevalence		Prevalence
BMI $\geq$ 30	0.386	Diagnosed with high blood pressure	0.663	Diagnosed with high blood pressure	0.753
Diagnosed with high blood pressure	0.373	Diagnosed with arthritis	0.473	Diagnosed with arthritis	0.684
Diagnosed with other serious chronic conditions	0.272	Diagnosed with other serious chronic conditions	0.433	Diagnosed with other serious chronic conditions	0.540
Smokes now	0.220	Has ever smoked cigarettes	0.437	Diff. heavy housework	0.492
Has ever smoked cigarettes	0.219	BMI $\geq$ 30	0.330	Diff. walking	0.402
Diagnosed with arthritis	0.179	Diagnosed with diabetes	0.280	Has ever smoked cigarettes	0.434
Diagnosed with psychological problems	0.119	Diagnosed with heart disease	0.190	Diagnosed with heart disease	0.374
Diagnosed with asthma	0.120	Diff. heavy housework	0.173	Diagnosed with cancer	0.342
Diagnosed with diabetes	0.110	Diagnosed with cancer	0.177	Diff. bathing/showering	0.176
Diff. heavy housework	0.075	Diff. walking	0.135	Diff. getting in/out of bed	0.180
Diff. walking	0.058	Diagnosed with psychological problems	0.136	Diff. getting outside	0.148
Diagnosed with lung disease	0.062	Diagnosed with asthma	0.140	Diagnosed with stroke	0.229
Diagnosed with heart disease	0.049	Diagnosed with lung disease	0.123	Diff. shopping for personal items	0.219
Diff. getting in/out of bed	0.036	Diagnosed with heart attack	0.120	Diagnosed with diabetes	0.219
Excessive alcohol drinking	0.044	Smokes now	0.117	Diff. dressing	0.157
Diagnosed with cancer	0.039	Diagnosed with stroke	0.082	Diagnosed with heart attack	0.236
Diagnosed with heart attack	0.028	Diff. getting in/out of bed	0.066	Diagnosed with loss of memory or mental ability	0.134
Diagnosed with loss of memory or mental ability	0.022	Diagnosed with loss of memory or mental ability	0.045	Diff. managing money	0.185
Diagnosed with stroke	0.024	Diff. bathing/showering	0.044	Diff. preparing meals	0.178
Diff. bathing/showering	0.018	Diff. light housework	0.043	Diff. using the phone	0.132
Diff. light housework	0.021	Diff. shopping for personal items	0.040	Diff. light housework	0.171
Diff. dressing	0.016	Diff. getting outside	0.033	Diff. using the toilet	0.104
Diff. shopping for personal items	0.018	Diff. dressing	0.033	Diagnosed with lung disease	0.180
Diff. getting outside	0.016	Diff. preparing meals	0.030	BMI $\geq$ 30	0.122
Diff. preparing meals	0.014	Diff. managing money	0.028	Diagnosed with psychological problems	0.122
Diff. managing money	0.014	Excessive alcohol drinking	0.028	Diagnosed with asthma	0.127
Diff. using the toilet	0.007	Diff. using the toilet	0.018	Diff. eating	0.048
Diff. eating	0.006	Diff. using the phone	0.015	Smokes now	0.035
Diff. using the phone	0.004	Diff. eating	0.010	Excessive alcohol drinking	0.002

Notes: Household heads between 45 and 49 (left); 65 and 69 (middle); 85 and 89 (right). PSID, waves 2005-2019.

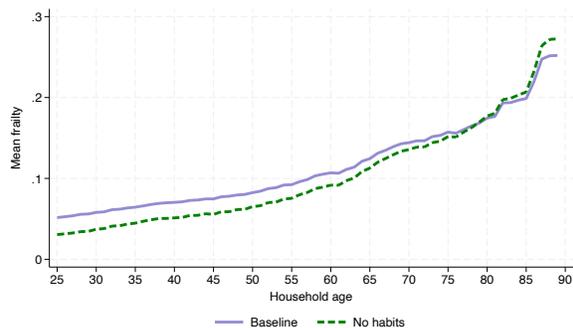
**Frailty excluding lifestyle habits.** The 29 deficits described in Table A-2 include four deficits related to obesity, smoking, and excessive alcohol consumption. Here, I examine whether excluding these deficits affects the distribution and dynamics of frailty, and the consumption response to a transitory health shock.

Excluding lifestyle-related deficits has only limited effects on the distribution of frailty over the life cycle. The alternative index has a lower overall mean (0.07 versus 0.09) and median (0.04 versus 0.07). Figure A-1a shows that average frailty by age is slightly lower under the alternative index, with the gap narrowing over the life cycle. Figure A-1b compares the 25th, 50th, and 75th percentiles by age and shows modest differences that are concentrated at younger ages. This pattern is consistent with Table 1, which documents that smoking and obesity are among the most prevalent deficits at younger ages. Figure A-1c indicates that the age profile of the variance is very similar across the two measures. Finally, Figure A-1d shows that the share of households with positive frailty is higher at younger ages under the alternative index. This reflects the fact that the excluded deficits are among the most prevalent conditions at younger ages. The difference narrows substantially after age 60.

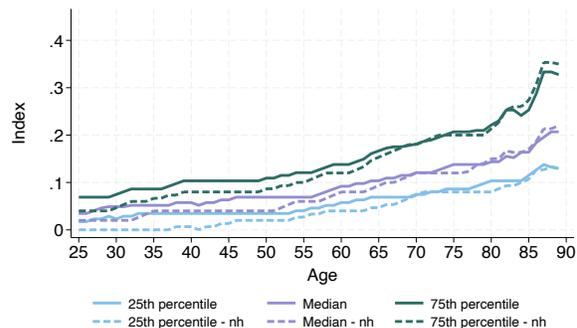
I next re-estimate the stochastic process for frailty using the alternative index. The estimation procedure is identical to that described in Section V.1, and full details are reported in Appendix C. Table A-4 shows that excluding lifestyle-related deficits has little effect on the estimated dynamics of frailty. The persistence parameter  $\rho_f$  increases slightly (0.91 versus 0.89), while the variance of the transitory shock ( $\sigma_{\varepsilon_f}^2$ ) rises, and the variance of the persistent shock ( $\sigma_{\eta_f}^2$ ) and the initial dispersion in the persistent component ( $\sigma_{\pi_0}^2$ ) decline modestly. In all cases, the differences are small relative to the reported standard errors. As a result, the implied persistence and volatility of health shocks are largely unchanged.

Finally, I re-estimate the pass-through coefficient to a transitory health shock using the two indices. The estimation procedure is identical to that described in Section V.2, and full details are reported in Appendix D. I estimate the pass-through coefficients jointly using a system GMM estimator, and Table A-5 reports the estimation results. This table suggests that excluding lifestyle habits has only limited effects on the consumption response to a transitory health shock. A clustered Wald test fails to reject equality of the two coefficients ( $p = 0.65$ ).

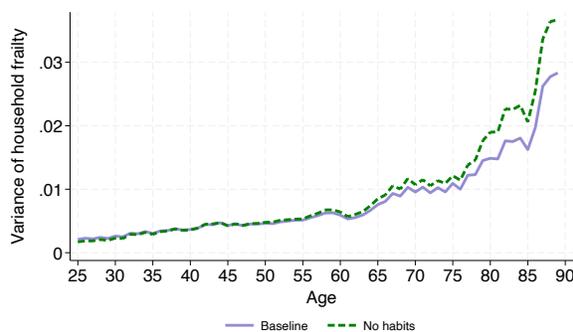
Figure A-1: Frailty statistics by age - Alternative indices



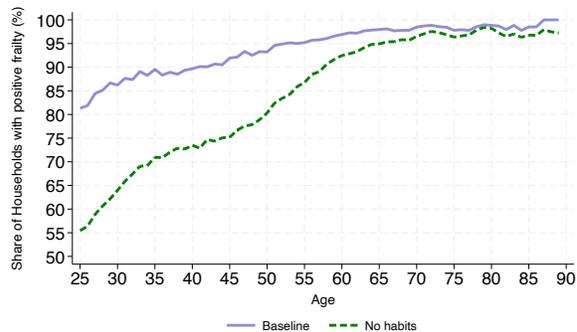
(a) Mean



(b) Percentiles



(c) Variance



(d) Share of households with positive frailty

Notes: Frailty statistics by age: Mean (a); 25th, 50th, 75th percentiles (b); variance (c); and share of households with positive frailty (d). Each statistic is smoothed using a 3-year moving average. PSID, 2005-2019.

Table A-4: Estimation results for non-zero frailty process

Parameter	Baseline	No lifestyle habits
$\rho_f$	0.89 (0.04)	0.91 (0.02)
$\sigma_{\varepsilon f}^2$	0.01 (0.02)	0.03 (0.01)
$\sigma_{\eta f}^2$	0.09 (0.03)	0.07 (0.02)
$\sigma_{\pi_0}^2$	0.34 (0.02)	0.30 (0.02)

Notes: The parameters of the stochastic component are estimated by GMM. PSID waves 2005-2019.

Table A-5: Pass-through of transitory health shocks to consumption

	Baseline frailty	No lifestyle habits
Transitory health shock ( $\phi_{\varepsilon f}$ )	-0.133** (0.054)	-0.153** (0.074)
N		10,158
$p$ -value: $\phi_{\varepsilon}^f = \phi_{\varepsilon,alt}^f$		0.652

Notes: Estimates obtained from a GMM procedure. Standard errors clustered at the household level. \* $p < 0.10$ , \*\* $p < 0.05$ , \*\*\* $p < 0.01$ .

## C First-Step Estimation

**Frailty process.** Table A-6 displays the estimation results for the probit regression for the probability of having zero frailty at each age. Figure A-2 displays the share of households with zero frailty in the raw data and as predicted by the probit regression.

Table A-6: Results from probit regression

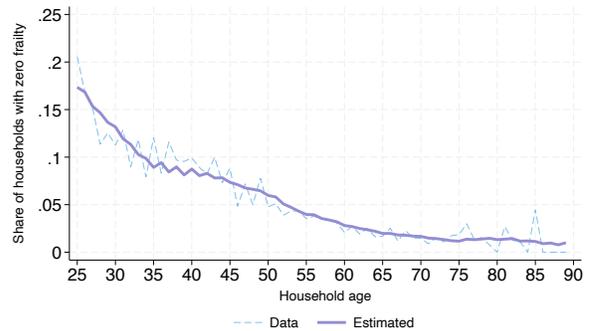
	Household has zero frailty
Age	-0.0574*** (0.0103)
Age <sup>2</sup>	0.000126 (0.000114)
Family size	-0.0427*** (0.00879)
Head's education	0.119*** (0.00572)
Constant	0.209 (0.256)
Cohort effects	Yes
Observations	28533
Pseudo $R^2$	0.0968

Standard errors in parentheses

\*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

Notes: Estimation results from zero frailty probit regression. The dependent variable equals 1 if the household has zero frailty.

Figure A-2: Share of households with zero frailty



Notes: Share of households with zero frailty in the data (dashed blue) and as predicted by the probit regression (solid purple).

After estimating the probability of zero frailty, I estimate the deterministic component of non-zero frailty ( $\kappa_t$  in Equation (3)) by regressing log-frailty on family size, education level, cohort effects, and a second-order polynomial in age. I then identify the parameters

of the stochastic process for non-zero frailty using the residuals from the regression for the deterministic component. In particular, I follow De Nardi, Fella, and Paz-Pardo (2019) and use a GMM procedure in which I minimize the distance between the theoretical and empirical age profiles of the variance and first-order autocovariance of residual frailty. Let  $\tilde{f}_{it} = \log f_{it} - \kappa_{it}$  denote residual frailty. The frailty process in Section V.1 implies that, for  $t > 1$ ,

$$\tilde{f}_{it} = \rho_f^{t-1} \pi_0^f + \sum_{j=2}^t \rho_f^{t-j} \eta_{ij}^f + \varepsilon_{it}^f,$$

from which I derive the following theoretical moments

$$\begin{aligned} \text{var}(\tilde{f}_{it}) &= \rho_f^{2(t-1)} \sigma_{\pi_0^f}^2 + \sum_{j=2}^t \rho_f^{2(t-j)} \sigma_{\eta^f}^2 + \sigma_{\varepsilon^f}^2, \\ \text{cov}(\tilde{f}_{it}, \tilde{f}_{i,t+1}) &= \rho_f^{2t-1} \sigma_{\pi_0^f}^2 + \sum_{j=2}^t \rho_f^{1+2(t-j)} \sigma_{\eta^f}^2. \end{aligned}$$

I compute the empirical age profile of the variance and first-order autocovariance of residual frailty as

$$\begin{aligned} \text{var}(\tilde{f}_{it}) &= \frac{1}{D} \sum_{d=1}^D \left( \frac{1}{N_{d,t}} \sum_{i=1}^{N_{d,t}} \tilde{f}_{it}^2 \right), \\ \text{cov}(\tilde{f}_{it}, \tilde{f}_{i,t+1}) &= \frac{1}{D} \sum_{d=1}^D \left( \frac{1}{N_{d,t}} \sum_{i=1}^{N_{d,t}} \tilde{f}_{it} \tilde{f}_{i,t+1} \right), \end{aligned}$$

where  $D$  is the number of waves in the dataset, and  $N_{d,t}$  is the number of observations of each age-wave cell.

I use a GMM procedure to minimize the distance between the theoretical and empirical moments using the identity matrix as the weighting matrix. I also compute standard errors by bootstrapping. Table A-7 reports the estimation results for the deterministic and stochastic components of frailty.

**Survival probabilities.** Table A-8 reports the estimation results for the logistic regression of a survival indicator for household heads.

Table A-7: Estimation results for non-zero frailty process

Log frailty			
Age	0.0259*** (0.00320)		
Age <sup>2</sup>	0.0000898*** (0.0000315)	$\rho_f$	0.89 (0.04)
Family size	-0.0638*** (0.00318)	$\sigma_{\varepsilon_f}^2$	0.01 (0.02)
Head's education	-0.0638*** (0.00184)		
Constant	-3.753*** (0.0867)	$\sigma_{\eta_f}^2$	0.09 (0.03)
Cohort effects	Yes	$\sigma_{\pi_0}^2$	0.34 (0.02)
Observations	26392		
$R^2$	0.215		

Standard errors in parentheses

\*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

Notes: Deterministic component (left) and parameters of the stochastic components (right.) The dependent variable for the deterministic component is log non-zero frailty. The parameters of the stochastic component are estimated by GMM. PSID waves 2005-2019.

Table A-8: Results for survival probabilities

Alive indicator	
Age	-0.0805 (0.0926)
Age <sup>2</sup>	0.000496 (0.000725)
Previous Period Frailty	-6.629*** (0.524)
Head's education	-0.00874 (0.0373)
Family size	0.361*** (0.104)
Cohort effects	Yes
Observations	25863
Pseudo $R^2$	0.176

Standard errors in parentheses

\*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

Notes: Estimation results for logistic regression of survival indicator. PSID waves 2005-2019.

**Earnings process.** The procedure for estimating earnings is similar to the one for estimating frailty described above. In particular, I first estimate the deterministic earnings profile ( $\kappa_t(f)$  in Equation (6)) by regressing the logarithm of earnings on frailty, family size, education level, cohort effects, and a second-order polynomial in age. I then use the residual earnings and a GMM procedure to estimate the four parameters governing the stochastic part of earnings.

Let  $\tilde{y}_t$  denote “detrended” log-earnings, that is, earnings net of the deterministic function  $\kappa(t, h)$ . Then, for  $t > 1$ , the theoretical moments are

$$\begin{aligned}\text{var}(\tilde{y}_{it}) &= \rho_y^{2(t-1)} \sigma_{\pi_0^y}^2 + \sum_{j=2}^t \rho_y^{2(t-j)} \sigma_{\eta^y}^2 + \sigma_{\varepsilon^y}^2, \\ \text{cov}(\tilde{y}_{it}, \tilde{y}_{i,t+1}) &= \rho_y^{2t-1} \sigma_{\pi_0^y}^2 + \sum_{j=2}^t \rho_y^{1+2(t-j)} \sigma_{\eta^y}^2.\end{aligned}$$

The corresponding empirical age profile of the variance and first-order autocovariance of residual earnings are

$$\begin{aligned}\text{var}(\tilde{y}_{it}) &= \frac{1}{D} \sum_{d=1}^D \left( \frac{1}{N_{d,t}} \sum_{i=1}^{N_{d,t}} \tilde{y}_{it}^2 \right), \\ \text{cov}(\tilde{y}_{it}, \tilde{y}_{i,t+1}) &= \frac{1}{D} \sum_{d=1}^D \left( \frac{1}{N_{d,t}} \sum_{i=1}^{N_{d,t}} \tilde{y}_{it} \tilde{y}_{i,t+1} \right),\end{aligned}$$

where  $D$  is the number of waves in the dataset, and  $N_{d,t}$  is the number of observations of each age-wave cell.

I use a GMM procedure to minimize the distance between the theoretical and empirical moments using the identity matrix as the weighting matrix. I also compute standard errors by bootstrapping. Table A-9 reports the estimation results for the deterministic profile of earnings and the parameters of the stochastic process.

**Out-of-pocket medical expenses.** Table A-10 reports the estimation results for the process for medical expenses. The variance of the i.i.d. shock to medical expenses is  $\sigma_{\xi}^2 = 0.043$

Table A-9: Estimation results for earnings process

Log household earnings			
Age	0.102*** (0.00579)		
Age <sup>2</sup>	-0.000953*** (0.0000670)	Parameter	Value
Household frailty	-3.431*** (0.0868)	$\rho_y$	0.99 (0.09)
Family size	0.110*** (0.00387)	$\sigma_{\varepsilon_y}^2$	0.09 (0.04)
Head's education	0.122*** (0.00247)	$\sigma_{\eta_y}^2$	0.01 (0.08)
Constant	6.824*** (0.114)	$\sigma_{\pi_0}^2$	0.36 (0.11)
Cohort effects	Yes		
Observations	24213		
$R^2$	0.239		

Standard errors in parentheses  
 \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

Notes: Deterministic component (left) and parameters of the stochastic components (right). The dependent variable for the deterministic component is log earnings. The parameters of the stochastic component are estimated by GMM. PSID waves 2005-2019.

Table A-10: Estimation results for medical expenses

	Log medical expenses	Squared Residuals
Age	0.155*** (0.00634)	-0.0988*** (0.0146)
Age <sup>2</sup>	-0.00127*** (0.0000632)	0.000974*** (0.000146)
Household frailty	0.394*** (0.120)	2.327*** (0.277)
Family size	0.217*** (0.00632)	0.0515*** (0.0146)
Head's education	0.130*** (0.00375)	-0.168*** (0.00866)
Constant	1.599*** (0.171)	4.850*** (0.396)
Cohort effects	Yes	Yes
Observations	28560	28560
$R^2$	0.147	0.0290

Standard errors in parentheses  
 \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

Notes: The first column is for the deterministic component. The second column is for the squared residuals from the regression in the first column. PSID waves 2005-2019.

## D Estimation of pass-through coefficients

**Detrending** I estimate pass-through coefficients for frailty and earnings by calculating the moments described in Section V.2 for detrended (i.e., net of deterministic components) log frailty, earnings, and consumption. As discussed in Commault (2022), the reason for using detrended values of such variables is to avoid mistaking as shocks (or as responses to shocks) the effect of demographic characteristics, such as age or family size. Following Blundell, Pistaferri, and Preston (2008) and Commault (2022), I detrend log frailty, earnings, and consumption by regressing them on age dummies, family size dummies, years of education, wave dummies, head’s race dummies, head’s employment status dummies, region of residence dummies, dummies for changes in family composition, dummies for the presence of dependent children living outside the family unit, and the interactions between wave dummies and years of education, race, employment status, and region of residence. I then estimate pass-through coefficients using the residuals from these regressions.

**Estimating restrictions** I estimate the pass-through of frailty and earnings shocks using the following estimating restrictions:

$$\begin{aligned} \mathbb{E} \left[ \Delta \log c_{it} \cdot \left( \tilde{\Delta} \log y_{i,t+1} \right) - \phi_{\varepsilon}^y \tilde{\Delta} \log y_{it} \left( \tilde{\Delta} \log y_{i,t+1} \right) \right] &= 0, \\ \mathbb{E} \left[ \Delta \log c_{it} \cdot \left( \tilde{\Delta} \log f_{i,t+1} \right) - \phi_{\varepsilon}^f \tilde{\Delta} \log f_{it} \left( \tilde{\Delta} \log f_{i,t+1} \right) \right] &= 0, \end{aligned}$$

Where  $\tilde{\Delta}$  denotes the quasi-difference  $\tilde{\Delta}x_{it} = x_{it} - \rho_x x_{i,t-1}$  for  $x = y, h$ ,  $\phi_{\varepsilon}^y$  denotes the pass-through coefficient for transitory earnings shocks, and  $\phi_{\varepsilon}^f$  denotes the pass-through coefficient for transitory frailty shocks.

**Estimation** I follow Commault (2022) and estimate the pass-through coefficients using the estimating restrictions above and a generalized method of moments. I pool all years together and estimate variances and covariances for the whole sample. Let  $X_i$  be the set of variables involved,  $\phi$  the vector of parameters, and  $g(X_i, \phi)$  the vector of estimating restrictions. The parameter estimates are the values that minimize a norm of the sample analog of the

moments:

$$\hat{\phi} = \underset{\phi_\varepsilon^y, \phi_\varepsilon^f}{\operatorname{argmin}} \left( \frac{1}{N} \sum_{n=1}^N g(X_n, \phi) \right)' \hat{W} \left( \frac{1}{N} \sum_{n=1}^N g(X_n, \phi) \right),$$

Where  $N$  is the number of household-year observations for which I observe the variables involved and  $\hat{W}$  is a weighting matrix. I choose  $\hat{W}$  so that the estimation of standard errors is robust to within-household correlations and heteroskedasticity.

**Estimated values** Table A-11 reports pass-through coefficients I estimate from my PSID sample. I find a positive consumption response to a transitory earnings shock. In particular, a 10% increase in earnings caused by a transitory earnings shock results in an increase of 1.7% in consumption. I also find a negative response to a transitory frailty shock. In this case, a 10% increase in frailty generates a 1.4% decrease in consumption.

Table A-11: Pass-through coefficients for transitory shocks

	All 25-61		All 25-89
$\phi_\varepsilon^y$	0.125***	$\phi_\varepsilon^f$	-0.103**
	0.02		0.05
N	11,095	N	12,444

**Results for steadily employed household heads.** As discussed in Section V.2, a potential concern is that transitory frailty shocks may affect consumption through labor force exit. To address this concern, I re-estimate the pass-through coefficient for households whose heads report being employed in periods three adjacent waves of the PSID ( $t - 1$ ,  $t$ , and  $t + 1$  for a generic  $t$ ). Table A-12 reports the results. These results show that the estimated pass-through remains negative when employment status is held fixed, although somewhat smaller in magnitude than in the full sample. This indicates that labor-force transitions are not the primary driver of the consumption response used to identify the effect of health on the marginal utility of consumption.

Table A-12: Frailty Pass-Through: Baseline and Stable Employment Sample

Sample	$\hat{\phi}_\varepsilon^f$	Std. Error	Observations
Full sample	-0.103**	0.05	12,444
Stable employment ( $t - 1, t, t + 1$ working)	-0.064	0.06	9,070

Notes: Estimates obtained from a GMM procedure. Standard errors clustered at the household level. \* $p < 0.10$ , \*\* $p < 0.05$ , \*\*\* $p < 0.01$ .

## E Details on the identification of $\delta$

In this section, I formalize the intuition for the identification of the effect of bad health on preferences. I follow Blundell et al. (2024), who use a similar argument but a different methodology.

The policy function for consumption provides information about the total effects of frailty shocks on consumption but is silent about the channels at work. In particular, the consumption policy function for workers in my model is:

$$c_t = c_t(a_t, \xi_t, \pi_t^y, \varepsilon_t^y, \pi_t^f, \varepsilon_t^f), \quad (\text{A1})$$

To analyze the channels at play, start with the Euler equation:

$$u_c(c_t, f_t) \geq s_{f,t} R \mathbb{E}[u_c(c_{t+1}, f_{t+1})], \quad (\text{A2})$$

where  $u(c_t, f_t) = (1 + \delta f_t) \frac{c_t^{1-\gamma}}{1-\gamma}$ ,  $u_c(\cdot)$  denotes the derivative of  $u(c_t, f_t)$  with respect to its first argument, and  $R = \beta(1 + r)$ . Then, using Equation (A1), the intertemporal budget constraint, and the laws of motion for  $\pi_t^y$  and  $\pi_t^f$ , rewrite the Euler equation as:

$$\begin{aligned} & u_c(c_t, f_t) \geq \\ & s_{f,t} R \\ & \mathbb{E}[u_c(c_{t+1}(a_t + y^n(ra_t + y_t(f)) + b_t - m_t(f) - c_t, \xi_{t+1}, \\ & \rho_y \pi_t^y + \eta_{t+1}^y, \varepsilon_{t+1}^y, \rho_f \pi_t^f + \eta_{t+1}^f, \varepsilon_{t+1}^f), \rho_f \pi_t^f + \eta_{t+1}^f + \varepsilon_{t+1}^f)] \end{aligned} \quad (\text{A3})$$

Equation (A3) relates current consumption  $c_t$  to the current state variables. It highlights the following channels at play:

1. Current frailty affects the marginal utility of current consumption—in purple,
2. Current frailty affects the survival probability—in green,
3. Assets, earnings, medical expenses, and government transfers affect the available resources after choosing current consumption. Available resources affect the next period's consumption and thus the value of current consumption that equalizes current and expected marginal utility—in blue,
4. The current persistent components of earnings and frailty affect the value of earnings and frailty in the next period and thus consumption in the next period—in orange

This optimality condition implicitly defines consumption as a function of these four channels. Thus, write log consumption as:

$$\log(c_t) = f\left( \underbrace{f_t}_{\text{MU}_c \text{ channel}}, \underbrace{f_t}_{\text{Survival channel}}, \underbrace{a_t + y^n(ra_t + y_t(f)) + b_t - m_t(f)}_{\text{Resource channel}}, \underbrace{\pi_t^y, \pi_t^f}_{\text{Future distributions channel}} \right), \quad (\text{A4})$$

Using Equation (A4), I analyze the consumption response to a transitory frailty shock. Because a transitory frailty shock does not affect the future distribution of frailty, it affects consumption only through the first three channels. Then, because people fully recover from a transitory shock within two years, I abstract from the effect of a transitory frailty shock on survival probabilities.<sup>10</sup> Thus, a transitory frailty shock affects consumption only through the marginal utility and resource channels.

The effect of a change in resources on consumption is the same regardless of whether the change is due to frailty or an earnings shock. As Blundell et al. (2024) notice, the effect on consumption (holding constant the ability to derive marginal utility from it) is the same whether people have to pay \$1,000 medical bill or earn \$1,000 less. Therefore, the effect

---

<sup>10</sup>Although the effect of transitory frailty shocks on survival is small, it may still cause fluctuations in consumption. I abstract from this for simplicity, but this effect is well-disciplined by the model.

of a change in resources is captured by the consumption response to a transitory earnings shock, which I measure with the pass-through coefficient  $\phi_e^y$ . This effect and the impact of a transitory frailty shock on medical expenses (which is known because I estimate medical expenses from the PSID and feed them into the model) give the hypothetical consumption response to a transitory frailty shock that would occur if only the resource channel were at play. Then, the effect of frailty on the marginal utility of consumption is identified residually from the overall pass-through coefficient to a transitory frailty shock and the one that would occur if only the resource channel were at play.

## F Computational details

**Solution.** The problem I describe in Section III has no analytical solution. Thus, I solve it numerically. I start from the final period of life (age 89) and proceed by backward iteration. I obtain policy functions for consumption and savings as functions of the household's state variables in each period. During the working years (ages 25 to 61), the state variables are age, assets, the shock to medical expenses, the persistent and transitory components of frailty, and the persistent and transitory components of earnings. During the retirement years (ages 63 to 89), the household's state variables include age, assets, the shock to medical expenses, and the persistent and transitory components of frailty. I discretize the endogenous and continuous variable for assets using a grid with 200 points. Then, I use the method in Rouwenhorst (1995) to discretize and approximate the stochastic processes for the persistent and transitory components of frailty and earnings and the shock to medical expenses using Markov chains. In particular, I discretize and approximate the AR(1) processes for  $\pi^y$  and  $\pi^f$  and the normally distributed shocks  $\xi$ ,  $\varepsilon^y$ , and  $\varepsilon^f$  using grids with 5 points each.<sup>11</sup> I obtain the asset policy function by optimizing the household's objective function using Brent's method. I compute the household's expected utility by integrating the value function over the distributions of the stochastic state variables. Using the intertemporal budget constraint and the asset policy function, I obtain the consumption policy function.

---

<sup>11</sup>Kopecky and Suen (2010) shows that the Rouwenhorst method with five grid points is more accurate than the Tauchen method with twenty-five.

**Simulation.** After obtaining the asset and consumption policy functions, I simulate the life cycle of 50,000 households. I initialize the simulations by drawing from the joint data distribution of frailty and wealth. Then, I simulate the household's frailty, earnings, and medical expenses using their laws of motion. Finally, based on the realizations of the state variables in each period, I simulate optimal consumption and savings starting at 25 and moving forward until 89 by interpolating the policy functions.

## G Pass-through coefficients of persistent shocks

The pass-through coefficients of persistent earnings and frailty shocks are defined as:

$$\phi_\eta^y = \frac{\text{cov}(\Delta \log c, \eta_t^y)}{\text{var}(\eta^y)}, \quad \phi_\eta^f = \frac{\text{cov}(\Delta \log c, \eta_t^f)}{\text{var}(\eta^f)}.$$

Following Kaplan and Violante (2010), I identify these pass-through coefficients using the following covariance restrictions:<sup>12</sup>

$$\begin{aligned} \text{cov}(\Delta \log c, \eta_t^y) &= \frac{1}{\rho_y} \text{cov}(\Delta \log c, \rho_y^2 \tilde{\Delta} \log y_{t-1} + \rho_y \tilde{\Delta} \log y_t + \tilde{\Delta} \log y_{t+1}), \\ \text{var}(\eta^y) &= \frac{1}{\rho_y} \text{cov}(\tilde{\Delta} \log y_t, \rho_y^2 \tilde{\Delta} \log y_{t-1} + \rho_y \tilde{\Delta} \log y_t + \tilde{\Delta} \log y_{t+1}), \\ \text{cov}(\Delta \log c, \eta_t^f) &= \frac{1}{\rho_f} \text{cov}(\Delta \log c, \rho_f^2 \tilde{\Delta} \log f_{t-1} + \rho_f \tilde{\Delta} \log f_t + \tilde{\Delta} \log f_{t+1}), \\ \text{var}(\eta^f) &= \frac{1}{\rho_f} \text{cov}(\tilde{\Delta} \log f_t, \rho_f^2 \tilde{\Delta} \log f_{t-1} + \rho_f \tilde{\Delta} \log f_t + \tilde{\Delta} \log f_{t+1}), \end{aligned}$$

where  $\tilde{\Delta}$  denotes the quasi-difference  $\tilde{\Delta}x_{it} = x_{it} - \rho_x x_{i,t-1}$  for  $x = y, f$ .

Estimating the pass-through coefficients of persistent shocks is similar to what I describe for transitory shocks in Appendix D. In particular, I use detrended log frailty, earnings, and consumption. I then estimate the pass-through coefficients using the same GMM procedure outlined in Appendix D and the following estimating restrictions:

$$\begin{aligned} \mathbb{E}[\Delta \log c \cdot (\rho_y^2 \tilde{\Delta} \log y_{t-1} + \rho_y \tilde{\Delta} \log y_t + \tilde{\Delta} \log y_{t+1}) - \phi_\eta^y \tilde{\Delta} \log y_t \cdot (\rho_y^2 \tilde{\Delta} \log y_{t-1} + \rho_y \tilde{\Delta} \log y_t + \tilde{\Delta} \log y_{t+1})] &= 0, \\ \mathbb{E}[\Delta \log c \cdot (\rho_f^2 \tilde{\Delta} \log f_{t-1} + \rho_f \tilde{\Delta} \log f_t + \tilde{\Delta} \log f_{t+1}) - \phi_\eta^f \tilde{\Delta} \log f_t \cdot (\rho_f^2 \tilde{\Delta} \log f_{t-1} + \rho_f \tilde{\Delta} \log f_t + \tilde{\Delta} \log f_{t+1})] &= 0. \end{aligned}$$

<sup>12</sup>These restrictions rely on the assumption that log consumption evolves as a random walk.

## References

- Achou, B. (2021). Housing liquidity and long-term care insurance demand: A quantitative evaluation. *Journal of Public Economics* 194, 104353.
- Achou, B., P. De Donder, F. Glenzer, M. Lee, and M.-L. Leroux (2023). At home versus in a nursing home: Long-term care settings and marginal utility. Working Paper 10482, CESifo.
- Ameriks, J., J. Briggs, A. Caplin, M. D. Shapiro, and C. Tonetti (2020). Long-term-care utility and late-in-life saving. *Journal of Political Economy* 128(6), 2375–2451.
- Arellano, M., R. Blundell, and S. Bonhomme (2017, May). Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework. *Econometrica* 85, 693–734.
- Attanasio, O. and S. J. Davis (1996). Relative wage movements and the distribution of consumption. *Journal of Political Economy* 104(6), 1227–1262.
- Bénabou, R. (2002). Tax and education policy in a heterogeneous-agent economy: What levels of redistribution maximize growth and efficiency? *Econometrica* 70(2), 481–517.
- Bernstein, S. L. and B. A. Toll (2019). Ask about smoking, not quitting: a chronic disease approach to assessing and treating tobacco use. *Addiction Science & Clinical Practice* 14(1), 1–4.
- Blundell, R., M. Borella, J. Commault, and M. De Nardi (2024, February). Old age risks, consumption, and insurance. *American Economic Review* 114(2), 575–613.
- Blundell, R., L. Pistaferri, and I. Preston (2008). Consumption inequality and partial insurance. *American Economic Review* 98(5), 1887–1921.
- Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2016, February). Consumption inequality and family labor supply. *American Economic Review* 106(2), 387–435.
- Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2018). Children, time allocation, and consumption insurance. *Journal of Political Economy* 126(S1), S73–S115.
- Bolt, U. (2021, November). What is the source of the health gradient? the case of obesity. Mimeo.
- Borella, M., F. Bullano, M. De Nardi, B. Krueger, and E. Manresa (2025, 05). Health inequality and health types. *The Econometrics Journal* 28(3), 341–384.

- Borella, M., M. De Nardi, M. Pak, N. Russo, and F. Yang (2023, 09). Fbbva lecture 2023. the importance of modeling income taxes over time: U.s. reforms and outcomes. *Journal of the European Economic Association* 21(6), 2237–2286.
- Braun, R. A., K. A. Kopecky, and T. Koreshkova (2016, 03). Old, Sick, Alone, and Poor: A Welfare Analysis of Old-Age Social Insurance Programmes. *The Review of Economic Studies* 84(2), 580–612.
- Brown, J. R., G. S. Goda, and K. McGarry (2016). Heterogeneity in state-dependent utility: Evidence from strategic surveys. *Economic Inquiry* 54(2), 847–861.
- Capatina, E. (2015). Life-cycle effects of health risk. *Journal of Monetary Economics* 74, 67 – 88.
- Carroll, C. D. (1997, 02). Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis. *The Quarterly Journal of Economics* 112(1), 1–55.
- CDC (2020). Hypertension prevalence among adults aged 18 and over: United states, 2017–2018. <https://www.cdc.gov/nchs/products/databriefs/db364.html>. Accessed April 2020.
- CDC (2021). Arthritis-related statistics. [https://www.cdc.gov/arthritis/data\\_statistics/arthritis-related-stats.html](https://www.cdc.gov/arthritis/data_statistics/arthritis-related-stats.html). Accessed October 2021.
- CDC (2022a). Adult obesity facts. <https://www.cdc.gov/obesity/data/adult.html>. Accessed August 2022.
- CDC (2022b). Burden of cigarette use in the u.s. <https://www.cdc.gov/tobacco/campaign/tips/resources/data/cigarette-smoking-in-united-states.html>. Accessed October 2022.
- Cochrane, J. H. (1991). A simple test of consumption insurance. *Journal of Political Economy* 99(5), 957–976.
- Commault, J. (2022, April). Does consumption respond to transitory shocks? reconciling natural experiments and semistructural methods. *American Economic Journal: Macroeconomics* 14(2), 96–122.
- Dal Bianco, C. (2023). Disability insurance and the effects of return-to-work policies. *Review of Economic Dynamics* 49, 351–373.
- Dal Bianco, C. and A. Moro (2022, October). The welfare effects of nonlinear health dynamics. Mimeo.
- De Nardi, M., G. Fella, and G. Paz-Pardo (2019, 03). Nonlinear Household Earnings Dynamics, Self-Insurance, and Welfare. *Journal of the European Economic Association* 18(2), 890–926.
- De Nardi, M., E. French, and J. B. Jones (2010). Why do the elderly save? the role of medical expenses. *Journal of Political Economy* 118(1), 39–75.

- De Nardi, M., E. French, and J. B. Jones (2016, November). Medicaid insurance in old age. *American Economic Review* 106(11), 3480–3520.
- De Nardi, M., S. Pashchenko, and P. Porapakarm (2017, October). The lifetime costs of bad health. Working Paper 23963, National Bureau of Economic Research.
- Deshpande, M. and L. M. Lockwood (2022). Beyond health: Nonhealth risk and the value of disability insurance. *Econometrica* 90(4), 1781–1810.
- Edwards, R. D. (2008). Health risk and portfolio choice. *Journal of Business & Economic Statistics* 26(4), 472–485.
- Evans, W. N. and W. K. Viscusi (1991). Estimation of state-dependent utility functions using survey data. *The Review of Economics and Statistics* 73(1), 94–104.
- Finkelstein, A., E. F. Luttmer, and M. J. Notowidigdo (2013). What good is wealth without health? The effect of health on the marginal utility of consumption. *Journal of the European Economic Association* 11(suppl.1), 221–258.
- French, E. (2005, 04). The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour. *The Review of Economic Studies* 72(2), 395–427.
- Gourinchas, P.-O. and J. A. Parker (2002). Consumption over the life cycle. *Econometrica* 70(1), 47–89.
- Gyrd-Hansen, D. (2017). A stated preference approach to assess whether health status impacts on marginal utility of consumption. *Health Economics* 26(10), 1224–1233.
- Haider, S. J. (2001). Earnings instability and earnings inequality of males in the united states: 1967–1991. *Journal of Labor Economics* 19(4), 799–836.
- Heathcote, J., K. Storesletten, and G. L. Violante (2014, July). Consumption and labor supply with partial insurance: An analytical framework. *American Economic Review* 104(7), 2075–2126.
- Heathcote, J., K. Storesletten, and G. L. Violante (2017). Optimal tax progressivity: An analytical framework. *The Quarterly Journal of Economics* 132(4), 1693–1754.
- Hong, J. H., J. Pijoan-Mas, and J.-V. Ríos-Rull (2015). Health heterogeneity and the preferences for consumption growth. Unpublished manuscript.
- Hosseini, R., K. A. Kopecky, and K. Zhao (2020). How Important Is Health Inequality for Lifetime Earnings Inequality? Working Paper.

- Hosseini, R., K. A. Kopecky, and K. Zhao (2022). The evolution of health over the life cycle. *Review of Economic Dynamics* 45, 237–263.
- Johnson, D. S., K. A. McGonagle, V. A. Freedman, and N. Sastry (2018). Fifty years of the panel study of income dynamics: Past, present, and future. *The ANNALS of the American Academy of Political and Social Science* 680(1), 9–28.
- Kaplan, G. and G. L. Violante (2010). How much consumption insurance beyond self-insurance? *American Economic Journal: Macroeconomics* 2(4), 53–87.
- Keane, M., E. Capatina, and S. Maruyama (2020, June). Health Shocks and the Evolution of Earnings over the Life-Cycle. Discussion Papers 2018-14b, School of Economics, The University of New South Wales.
- Koijen, R. S., S. Van Nieuwerburgh, and M. Yogo (2016). Health and mortality delta: Assessing the welfare cost of household insurance choice. *The Journal of Finance* 71(2), 957–1010.
- Kools, L. and M. Knoef (2019). Health and consumption preferences; estimating the health state dependence of utility using equivalence scales. *European Economic Review* 113, 46–62.
- Kopecky, K. A. and T. Koreshkova (2014, July). The impact of medical and nursing home expenses on savings. *American Economic Journal: Macroeconomics* 6(3), 29–72.
- Kopecky, K. A. and R. M. Suen (2010). Finite state markov-chain approximations to highly persistent processes. *Review of Economic Dynamics* 13(3), 701–714.
- Kydland, F. E. and E. C. Prescott (1982). Time to build and aggregate fluctuations. *Econometrica* 50(6), 1345–1370.
- Lillard, L. A. and Y. Weiss (1997). Uncertain health and survival: Effects on end-of-life consumption. *Journal of Business & Economic Statistics* 15(2), 254–268.
- McGee, R. (2021, December). Old age savings and house price shocks. Mimeo.
- Michaud, P.-C. and P. St. Amour (2023). Longevity, health and housing risks management in retirement. Technical report, National Bureau of Economic Research.
- Nygaard, V. (2021, May). Causes and consequences of life expectancy inequality. Mimeo.
- Palumbo, M. G. (1999, 04). Uncertain Medical Expenses and Precautionary Saving Near the End of the Life Cycle. *The Review of Economic Studies* 66(2), 395–421.
- Paz-Pardo, G. (2022, April). Homeownership and portfolio choice over the generations. Mimeo.

- Rouwenhorst, K. (1995). Asset pricing implications of equilibrium business cycle models. In P. U. Press (Ed.), *Frontiers of Business Cycle Research*.
- Russo, N., R. McGee, M. De Nardi, M. Borella, and R. Abram (2024). Health inequality and economic disparities by race, ethnicity, and gender. Working Paper 32971, National Bureau of Economic Research.
- Rust, J. and C. Phelan (1997). How social security and medicare affect retirement behavior in a world of incomplete markets. *Econometrica* 65(4), 781–831.
- Salvati, F. (2020). Health inequality, labor supply and retirement policies. Mimeo.
- Scholz, J. K. and A. Seshadri (2011, May). Health and wealth in a life cycle model. Research Paper 2010-224, Michigan Retirement Research Center. Available at SSRN: <https://ssrn.com/abstract=1706918>.
- Scholz, J. K., A. Seshadri, and S. Khitatrakun (2006). Are americans saving “optimally” for retirement? *Journal of Political Economy* 114(4), 607–643.
- Searle, S. D., A. Mitnitski, E. A. Gahbauer, T. M. Gill, and K. Rockwood (2008). A standard procedure for creating a frailty index. *BMC geriatrics* 8(1), 1–10.
- Sloan, F. A., W. Kip Viscusi, H. W. Chesson, C. J. Conover, and K. Whetten-Goldstein (1998). Alternative approaches to valuing intangible health losses: the evidence for multiple sclerosis<sup>1</sup>this article was written when chesson was a student at duke university.<sup>1</sup> *Journal of Health Economics* 17(4), 475–497.
- Viscusi, W. K. and W. N. Evans (1990). Utility functions that depend on health status: Estimates and economic implications. *The American Economic Review* 80(3), 353–374.
- Voena, A. (2015, August). Yours, mine, and ours: Do divorce laws affect the intertemporal behavior of married couples? *American Economic Review* 105(8), 2295–2332.
- Wu, C. and D. Krueger (2018, March). How much consumption insurance in bewley models with endogenous family labor supply? Working Paper 24472, National Bureau of Economic Research.