The Costs of Bad Mental and Physical Health

Joaquín Galeno Ghio

University of Minnesota

November 2025

Abstract

How do mental and physical health jointly shape individuals' economic outcomes and saving behavior over the life cycle? Using data from the PSID and MEPS, I document that (i) mental health problems peak early in life and improve through mid-adulthood, while physical health steadily declines, (ii) poor mental and physical health each lowers earnings and raise medical spending, but their joint occurrence leads to the largest losses, and (iii) high-productivity individuals enjoy better health transitions and lower mortality, linking economic opportunity to health. I embed these facts in a structural life-cycle model that separates mental and physical health risks and allows preferences, productivity, and medical expenses to interact. The estimation shows that people value consumption more when in poor health, which alters their saving behavior. Health shocks reduce lifetime resources by about 3% of average earnings—twice as much for low-productivity individuals. Counterfactuals reveal that worse early-life health mainly shifts consumption toward youth, tightening midlife savings. Together, the results show that multidimensional health risks are closely linked to differences in wealth accumulation and lifetime welfare.

1 Introduction

Health shapes people's economic lives. Poor health reduces productivity and labor market participation, and it alters how households save and consume over the life cycle French (2005); De Nardi et al. (2024); Hosseini et al. (2022); Kools and Knoef (2019); Russo (2023). Yet most economic research has treated health as a single dimension—typically associated with their physical condition—overlooking mental health, despite its major role in determining earnings, labor supply, and longevity.

The neglect of mental health is increasingly costly. Across many countries, younger cohorts report sharp declines in psychological well-being Blanchflower and Bryson (2024); Blanchflower et al. (2024a,b). This deterioration raises concerns about future earnings inequality, social insurance systems, and the capacity of households to build wealth. Understanding how mental and physical health jointly shape economic behavior is therefore crucial for both research and policy.

Recent work has begun to tackle this question. Abramson et al. (2025) develop a theory of mental illness that links cognitive frictions and stigma to macroeconomic outcomes—consumption, savings, portfolio choice, and labor supply. Ascarza-Mendoza et al. (2024) documents the life-cycle dynamics of mental health using U.S. data, showing sizable welfare losses from depressive symptoms and proposing a parsimonious statistical model for structural applications. These studies establish mental health as an economically meaningful source of risk.

I build on this literature by being the first work to integrate mental and physical health together in a structural life-cycle model. The framework links health shocks, earnings, and saving decisions, allowing me to isolate how each health dimension—and their interaction—shapes economic behavior. This unified treatment reveals new mechanisms behind the income and welfare changes associated with poor health.

Using microdata from the Panel Study of Income Dynamics (PSID) and medical expenditure data from the Medical Expenditure Panel Survey (MEPS), I document distinct life-cycle patterns of mental and physical health. Mental health problems are most common in early life and tend to improve through mid-adulthood, whereas physical health declines steadily with age. Medical spending rises sharply with age, and differences across health states are largest before age 60. These facts motivate modeling medical expenses explicitly, capturing both the likelihood of incurring costs and the amount spent when health shocks occur.

Earnings data analysis reveals that poor mental and physical health each reduce income, but their joint occurrence generates the most considerable losses. I also uncover a strong two-way link between productivity and health: individuals with higher earnings potential experience better health transitions and lower mortality rates. To capture this interaction, I introduce productivity types into the model, allowing health dynamics to depend on productivity. This feature ties economic opportunity and health trajectories together in a unified framework.

Building on these findings, I estimate a structural life-cycle model of consumption and saving with productivity heterogeneity, medical expenses, health-dependent preferences, bequest motives, and exogenous retirement. In a first step, I estimate the empirical processes for earnings, medical spending, and health dynamics. In a second step, I use the Simulated Method of Moments to identify preference parameters, bequest motives, and the government-provided consumption floor from age–wealth profiles by health status. The results show that individuals value consumption more when in poor health, a feature that has significant implications for saving incentives, precautionary motives, and the role of social insurance.

I then use the model to study the macroeconomic implications of worsening early-life health. In a counterfactual where the initial distribution of mental health at age

25 deteriorates, individuals consume more early in life and save less, even though their long-run income and wealth paths remain essentially unchanged. These results suggest that policies improving mental health at young ages can shift the timing of consumption and welfare, even without large effects on lifetime earnings or asset accumulation.

Motivated by recent evidence showing rising chronic conditions and worsening physical health among younger cohorts, I conduct a second counterfactual that deteriorates the initial physical health distribution at age 25. A larger share of individuals starting life in poor physical health generates stronger income effects, since earnings fall sharply when bad mental and physical health coincide. As a result, the consumption is almost unchanged, but wealth accumulation drops steeply: early-life asset holdings decline by nearly 20 percent. These findings suggest that early declines in physical health can have long-lasting consequences for wealth building, even when short-run consumption effects are limited.

Finally, I quantify the monetary costs of poor health over the life cycle by comparing each individual with a counterfactual self who never experiences bad mental or physical health. This exercise measures the combined income and medical spending losses associated with health shocks across different productivity types. The results show that eliminating mental health shocks reduces lifetime losses by about 1.1 percent of average earnings, while removing physical health shocks yields similar relief (1.3 percent). When both types of shocks are jointly suppressed, total losses reach 3.2 percent of average earnings, revealing strong complementarity between mental and physical health risks. These costs are not evenly distributed: low-productivity individuals bear losses equivalent to 5.3 percent of their average earnings, compared to 2.5 percent among high-productivity workers. Poor health, therefore, imposes a disproportionate financial burden on the least productive, highlighting the welfare gains from policies that mitigate health shocks.

The paper proceeds as follows. Section 2 reviews the related literature. Section 3 describes the data and health measures. Section 4 presents the model and Section 5 the estimation strategy. Section 6 examines health-dependent preferences, Section 7 reports counterfactuals, and Section 8 quantifies the monetary costs of poor health. Section 9 concludes.

2 Related Literature

This paper connects to three strands of research and contributes to each. First, it relates to the literature on structural life-cycle models with health risk. Seminal contributions include Palumbo (1999), French (2005), and De Nardi et al. (2010), followed by Capatina (2015), Capatina et al. (2020), Capatina and Keane (2025), De Nardi et al. (2024), and Hosseini et al. (2025). These studies model health risk through medical expenditures, survival, and insurance institutions, analyzing their implications for saving, labor supply, and welfare. I build on this work by being the first to explicitly distinguish between mental and physical health within a structural life-cycle framework. Each dimension exhibits its own persistence, earnings consequences, and effects on preferences. This distinction isolates mechanisms that are otherwise conflated in one-dimensional health indexes and clarifies how different types of health risk shape consumption and saving decisions over the life cycle.

Second, the paper contributes to the emerging economics of mental health. Abramson et al. (2025) develops a life-cycle model grounded in psychiatric evidence to map how mental health affects economic choices. Jolivet and Postel-Vinay (2024) analyzes the joint dynamics of mental health and labor-market outcomes in a search framework. On the empirical side, Ascarza-Mendoza et al. (2024) documents the transition patterns of mental health in the U.S. and proposes a statistical model suitable for structural applications. I extend this line by jointly documenting the dynamics of mental and physical health and embedding these processes in a structural model that tracks both dimensions simultaneously. The model accounts for cross-state disparities and is used to evaluate counterfactuals—such as worsening mental health at entry—that are hard to study empirically.

Third, the paper relates to research on health-dependent preferences. As emphasized by Russo (2023), there is little consensus on whether poor health raises or lowers the marginal utility of consumption. Early work, such as Palumbo (1999), finds limited effects, while De Nardi et al. (2010) and Russo (2023) report that marginal utility falls as health worsens. Other studies—including Lillard and Weiss (1997), Kools and Knoef (2019), and Ameriks et al. (2020)—find the opposite: that poor health increases the value of consumption. I contribute by quantifying health-dependent preference shifters in both mental and physical dimensions within a unified life-cycle framework, and by assessing their implications for consumption, saving, and bequests.

Taken together, the paper integrates multidimensional health, mental health dynamics, and health-dependent preferences into a single structural framework, providing a coherent explanation for how health risks influence saving behavior, income profiles, and welfare over the life cycle.

3 Data

The empirical analysis draws primarily on the Panel Study of Income Dynamics (PSID), using waves from 2001–2003 and 2007–2021.[1] Over this period, the PSID is a biennial, nationally representative survey of U.S. households conducted by the University of Michigan's Survey Research Center. It provides rich longitudinal data on economic, demographic, and health outcomes, making it ideal for studying the interaction between health and economic behavior. I restrict the sample to household heads aged

[1] The 2005 wave is excluded because it lacks information on respondents' mental health.

25–90 who respond to the survey.

To incorporate medical expenditures, I complement the PSID with the Medical Expenditure Panel Survey (MEPS), a rotating panel covering 2000–2022. MEPS offers detailed information on both medical spending and health conditions, including out-of-pocket payments and self-reported health status. Crucially, it uses the same health questions as the PSID, ensuring direct comparability between the two datasets.

Income is measured at the household level as the sum of labor, farm, and business earnings. Wealth corresponds to net worth, including checking and savings accounts, business and real estate assets, stocks, vehicles, private annuities or IRAs, and home equity. All monetary variables are expressed in 2010 U.S. dollars using the Consumer Price Index (CPI) from the Bureau of Labor Statistics, while medical expenditures are deflated using the Personal Consumption Expenditures Health (PCE-Health) index.

3.1 Physical and Mental Health

Physical Health: Physical health is measured using the self-reported health status available in the PSID. In each wave, respondents are asked: *Would you say your health in general is excellent, very good, good, fair, or poor?* Following the standard practice in the literature (French (2005), De Nardi et al. (2024), Mahler and Yum (2024)), I define a "bad physical health" indicator equal to one if the respondent reports "fair" or "poor" health, and zero otherwise. This dummy captures the lower tail of the health distribution and is widely used as a reliable proxy for poor physical condition in both structural and empirical studies.

Mental Health: Mental health is measured using the Kessler Psychological Distress Scale (K-6), a short an validated screening instrument for psychological distress. The K-6 has been shown to perform well in identifying moderate and severe mental illness in mutiple settings—within the U.S. Kessler et al. (2003); Prochaska et al. (2012);

Kessler et al. (2002); Umucu et al. (2022) and abroad Cornelius et al. (2013); Furukawa et al. (2003); Kessler et al. (2010). I classify individuals into two mental health states, Good or Bad, based on the K-6 thresholds for moderate distress defined by Prochaska et al. (2012).

Sample selection: Because the mental health measure is collected only for respondents, while the physical-health question is asked of both the household reference person and their spouse, I restrict the analysis to reference persons who are also PSID respondents. To ensure representativeness and comparability with the broader U.S. population and existing studies, I exclude the Survey of Economic Opportunity (SEO) subsample, which oversampled low-income and minority households. Both men and women are included in the analysis.

Finally, because the PSID is collected on a biennial basis during my study period, I group individuals into two-year age bins. This choice ensures that the empirical moments are directly comparable to the model and helps smooth out short-term fluctuations that could otherwise distort the analysis. A more detailed discussion of the data structure and the sample selection process is provided in Appendix A.

The Dynamics of Health: Figure 1 shows the share of individuals in bad mental health and bad physical health by age. The two series exhibit markedly different life-cycle patterns. Mental health problems are most prevalent in early adulthood and decline steadily through midlife, suggesting that psychological distress eases as individuals age. Around age 75, this trend reverses slightly, with mental distress becoming more common at older ages. Physical health follows the opposite trajectory: the share of individuals in poor physical condition rises monotonically with age, reflecting the gradual but persistent deterioration of the body.

These patterns align with earlier evidence. The steady increase in poor physical health is consistent with Mahler and Yum (2024) and De Nardi et al. (2024), while the improving profile of mental health through early and mid-adulthood matches the findings of Ascarza-Mendoza et al. (2024) for individuals under 70.

Figure 1: Share of individuals in bad Mental and Physical Health by age

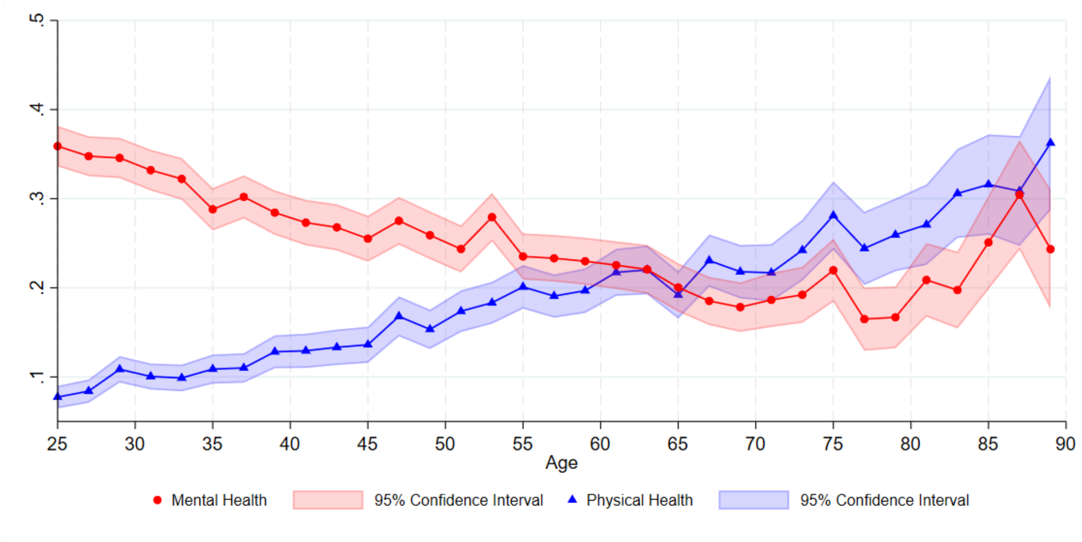

Beyond confirming earlier results, the figure also illustrates why it is essential to treat mental and physical health as separate dimensions. Mental health shocks are concentrated among younger adults, while physical health deteriorates gradually and overtakes mental distress around age 60. Collapsing these dynamics into a single health index hides the timing and nature of the risks individuals face. Distinguishing between the two allows the analysis to capture how each dimension shapes economic decisions and well-being at different stages of the life cycle.

Medical Expenses: Health can affect medical spending through two channels. On the extensive margin, poor health increases the likelihood of incurring in any medical expenses. On the intensive margin, conditional on spending, health shocks raise the amount individuals pay for health good and services.

Figure 2 illustrates both margins using MEPS data. Panel 2a shows the share of individuals with positive out-of-pocket medical expenditures[2] by age. As expected, this share rises steadily with age. Differences across health states are largest before age 60: individuals in good mental and physical health are about 10 percentage points less likely to incur medical expenses than those in poor health in both dimensions. After age 60, these gaps narrow, as medical spending becomes nearly universal across health groups.

Figure 2: Medical Expenses by age

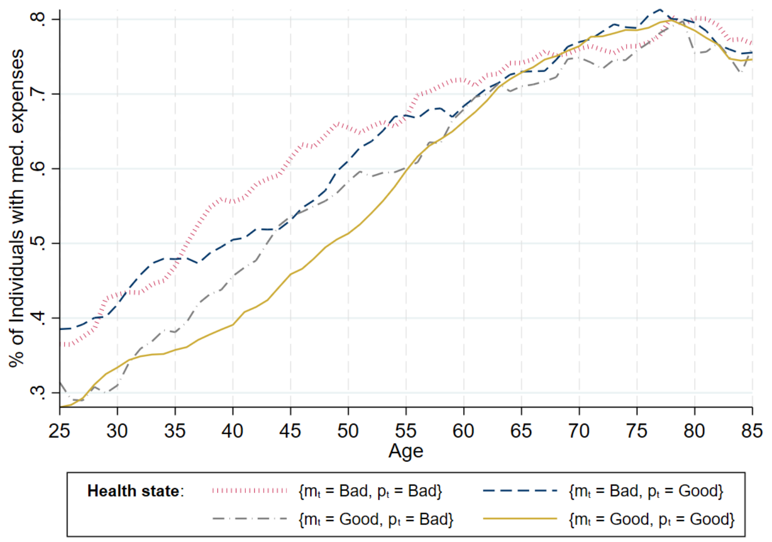

(a) Share of individuals with positive medical expenses by age

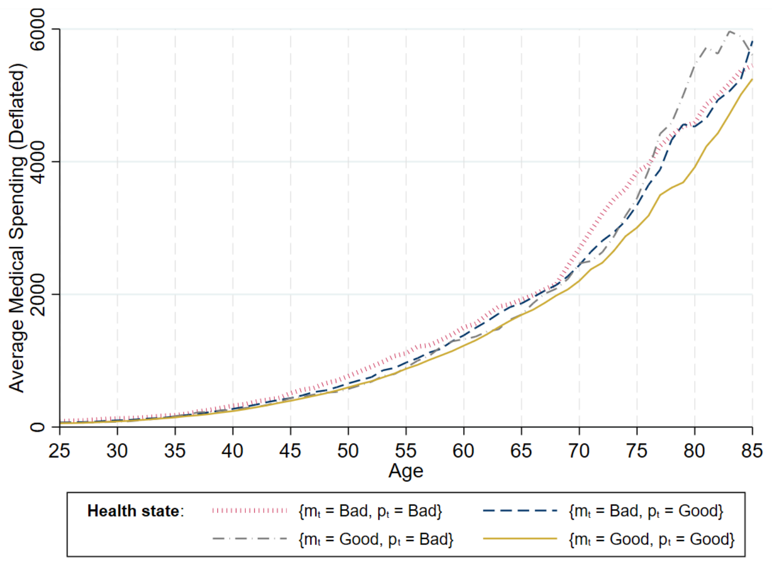

(b) Average amount of medical expenses by age

[2]Because spending amounts are noisy, I classify individuals as having medical expenses only if their annual spending exceeds \$120, equivalent to \$10 per month.

On the intensive margin, Figure 2b plots the average level of out-of-pocket medical expenses net of individual effects.[3] The average amount of spending rises steadily with age but shows no systematic differences across health states, suggesting that once individuals incur medical expenses, the amount paid is largely driven by age rather than health status.

4 Model

This section presents a dynamic life-cycle model of consumption and saving in a risk-free asset under exogenous health shocks. The framework features exogenous retirement, bequest motives, heterogeneous productivity types, medical expenses, health-dependent preferences, a consumption floor, and —crucially— the explicit distinction between mental and physical health as separate dimensions of the health process.

Time is discrete and indexed by t, with each period representing two years. The model considers a household head seeking to maximize their expected lifetime utility starting at age 25. Accordingly, the model spans ages $t = t_{25}, t_{27}, t_{29}, \cdots, T$. Individuals retire exogenously at age $T_r = 65$, and their lives extend up to a maximum age of $T = 89$. Each period, they face four key sources of uncertainty: health shocks, mortality risk, medical expenses risk, and earnings shocks. Given these risks, individuals choose consumption and saving in a risk-free savings asset.

Individuals begin life with assets equal to the median wealth in the PSID for their initial health state. They can then save but not borrow, and the risk-free asset yields a constant return r.

Preferences: In each period, utility depends on consumption c_t and current health state h_t. Preferences are time-separable, and future utility is discounted by a constant

[3] To isolate variation unrelated to persistent individual characteristics, I use residual medical expenses after controlling for individual fixed effects.

discount factor β. The contemporaneous utility function is given by:

$$u(c_t, h_t) = \delta(p_t)\psi(m_t)\frac{c_t^{1-\gamma}}{1-\gamma} \tag{1}$$

Where $\delta(p_t)$ shifts marginal utility according to the individual's physical health status, and $\psi(m_t)$ does so according to mental health. Consumption is equivalized by the square root of household size to account for scale economies.[4]

Following De Nardi et al. (2010), both health dimensions enter multiplicatively as binary shifters

$$\delta(p_t) = 1 + \delta\mathbb{1}\{p_t = \text{bad}\} \tag{2}$$

$$\psi(m_t) = 1 + \psi\mathbb{1}\{m_t = \text{bad}\} \tag{3}$$

Unobserved heterogeneity and productivity types: Each individual belongs to a fixed productivity type that remains constant over the life cycle. Following Low and Pistaferri (2015), I estimate individual fixed effects in the earnings process, denoted by κ, and classify individuals into three groups: the bottom quartile (low productivity, z_L), the middle 50 percent (medium productivity, z_M), and the top quartile (high productivity, z_H). This classification captures long-term earnings potential in a parsimonious manner.

Although κ is continuous in the data, I discretize it into ten support points in the model for tractability. These values are drawn from a normal distribution whose mean and standard deviation match the empirical distribution of κ, preserving observed variability while keeping the model computationally manageable.

Productivity type also shapes health dynamics. Health transition probabilities depend on z, meaning that individuals with lower productivity not only earn less but also face

[4]This equivalization follows standard practice in the literature.

a higher risk of health deterioration. This interaction strengthens the link between economic opportunities and health outcomes, highlighting the dual role of productivity in driving both income and health inequalities over the life cycle.

Health: In each period, individuals face uncertainty about their health state, represented as a two-dimensional vector, $h_t = (m_t, p_t)$, where m_t denotes mental health and p_t represents physical health.

Each dimension can take two states, *good* or *bad*, plus an absorbing *death* state. Health evolves according to a first-order Markov process conditional on age t and productivity type z :

$$\pi^{h'}_{h,z,t} = Pr(h_{t+1} = k | h_t, z, t) \tag{4}$$

Mortality enters through age-dependent survival probabilities. Let ϕ_t denote the unconditional probability of being alive at age t, with $\phi_{25} = 1$. The conditional survival probability of reaching age t, given survival up to $t-1$, is then $\zeta_t = \frac{\phi_t}{\phi_{t-1}}$. Death is modeled as an absorbing state in the Markov chain.

Individuals enter the model at age 25 with an initial distribution of health states that matches the empirical distribution observed in the PSID, ensuring that simulated health trajectories align with observed data.

Medical expenses: As documented in Section 3, medical expenditures exhibit two salient patterns. At younger ages, many individuals report no out-of-pocket spending, while at older ages almost everyone incurs positive expenses. Conditional on spending, the average amount rises sharply with age.

To capture these features, I model medical expenses in two steps. First, the probability of incurring any out-of-pocket expenses depends on age and health:

$$\pi_{h,t}^{me} = Pr(me_t | h_t, t) \tag{5}$$

Second, conditional on having positive expenses, I follow Yu (2024) and model the level of spending as a function of age and health, $e(h_t, t)$. This two-step structure enables the model to match both the extensive margin—the likelihood of spending—and the intensive margin—the amount spent—over the life cycle. The overall process for medical expenses can be written as

$$e_t = e(me_t, h_t, t) = \begin{cases} e(h_t, t) & \text{if } me_t = 1 \\ 0 & \text{if } me_t = 0 \end{cases}$$

Earnings Process: The life-cycle is divided into two stages: the working period and retirement. During working years, labor supply is inelastic, and households face earnings risk. Earnings consist of a deterministic component and a stochastic component. The deterministic part depends on age, mental health, physical health, the individual fixed effect κ_i, and the interaction between the health variables.

The stochastic component captures persistent idiosyncratic shocks. The persistent term follows a discretized AR(1) process with persistence ρ and innovation variance σ_v^2, approximated using the method of Rouwenhorst (1995). Formally, the earnings process corresponds to:

$$\log y_{i,t} = \kappa_i + \alpha(age_t) + \theta \mathbb{1}\{m_t = bad\} + \gamma \mathbb{1}\{p_t = bad\} + \mu \mathbb{1}\{m_t = bad \wedge p_t = bad\} + \log \varepsilon_t \tag{6}$$

Where κ_i denotes the individual fixed effect, $\alpha(age_t)$ is a quadratic polynomial capturing the common age profile of earnings. The coefficients θ, γ, and μ quantify the earnings penalties associated with being in bad mental health, bad physical health, or both simultaneously, respectively. The stochastic component $\log \varepsilon_t$ evolves as an AR(1):

$$\log \varepsilon_t = \rho \log \varepsilon_{t-1} + \eta_t \quad \text{with} \quad \eta_t \sim N(0, \sigma_\eta^2) \tag{7}$$

At age 65, individuals retire. Following De Nardi et al. (2019), post-retirement earnings are modeled as a deterministic function of the last earnings realization from before retirement:

$$y^{ret} = y(y_{T^{ret}-1}) \tag{8}$$

This formulation reflects the strong empirical link between retirement income and late-career earnings. The exact functional form used to approximate retirement earnings is described in Appendix B.

Bequest motives: As in De Nardi (2004), individuals derive utility not only from consumption while alive but also from the assets they leave at death. Let a denote bequeathed assets. The utility from bequest is given by:

$$\Theta(a) = \theta_{beq} \frac{(a + \kappa_{beq})^{1-\sigma}}{1-\sigma}$$

Where θ_{beq} captures the intensity of the bequest motive, and κ_{beq} determines the extent to which bequests are a luxury good. A higher κ_{beq} implies that only wealthier households place significant value on leaving bequests.

Government: The government provides transfers, b_t, to guarantee a minimum consumption floor $\underline{c}$. Transfers covers the gap between available resources and the sum of required consumption and medical expenses:

$$b_t = \max\{0, \underline{c} + e(h_t, t) - (a_t + y_{i,t}(h_t, \kappa, \varepsilon_t, t))\} \qquad \text{if } t < 65 \tag{9}$$

$$b_t = \max\{0, \underline{c} + e(h_t, t) - (a_t + y_{i,t}(y_{63}))\} \qquad \text{if } t \geqslant 65 \tag{10}$$

Timing: Each period begins with the household holding assets and a fixed productivity type. During working years, individuals first draw realizations of earnings, health, and medical expense shocks, and then choose consumption and savings. The medical expense shock determines whether the individual incurs out-of-pocket spending; if so, the amount depends on age and health status.

After retirement, earnings are predetermined, but health and medical expense shocks continue to arrive. At the end of each period, a survival shock determines whether the individual continues to the next period or leaves a bequest. Survivors transition with updated assets and health states, while those who die derive utility from bequests. This timing follows closely the structure in De Nardi et al. (2010).

Individual's Decision Problem: The household's optimization problem is characterized by two value functions—one for working-age individuals and another for retirees.

Working-age individuals $t < 65$: For individuals who have not yet retired, the state vector is $X_t = \{t, a_t, me_t, \kappa, h_t, \varepsilon_t\}$, where t denotes age, a_t current assets, me_t a binary indicator for incurring out-of-pocket medical expenses. The term κ represents the individual fixed productivity effect[5], $h_t = (m_t, p_t)$ captures the two-dimensional health state (mental and physical), and ε_t is the persistent productivity shock. Individuals solve the following problem:

$$V_t(X_t) = \max_{c_t, a_{t+1}} \{u(c_t, h_t) + \beta \{\zeta_{t+1,z}(h_{t+1})\mathbf{E}_t [V_{t+1}(X_{t+1})] + (1 - \zeta_{t+1,z}(h_{t+1}))\Theta(a_{t+1})\}\}$$

[5]The individual fixed effect determines both an individual's productivity level and the productivity type to which they belong.

This problem is subject to the standard intertemporal budget and borrowing constraints:

$$c_t + a_{t+1} \leqslant (1+r)a_t + y_t(h_t, \kappa, \varepsilon_t) - e_t(me_t, h_t, t) \tag{11}$$

$$a_{t+1} \geqslant 0 \tag{12}$$

Along with the laws of motion for health and earnings described in equations (4)-(7).

Retired individuals, $t \geqslant 65$: For retirees, the state vector is $X_t = \{t, a_t, me_t, z, h_t, y_{63}\}$, where z denotes the productivity type, and y_{63} the last labor income realization before retirement, which determines post-retirement earnings. Retirees maximize:

$$V_t(X_t) = \max_{c_t, a_{t+1}} \{u(c_t, h_t) + \beta \{\zeta_{t+1,z}(h_{t+1})\mathbf{E}_t [V_{t+1}(X_{t+1})] + (1 - \zeta_{t+1,z}(h_{t+1}))\Theta(a_{t+1})\}\}$$

Because retirees no longer face earnings shocks, their simplifies to:

$$c_t + a_{t+1} \leqslant (1+r)a_t + y^{ret}(y_{63}) - e(me_t, h_t, t) \tag{13}$$

$$a_{t+1} \geqslant 0 \tag{14}$$

Together with the health transition process (4) and the retirement income specification (8).

5 Estimation

This section describes the estimation strategy, reports parameter estimates, and assesses the model's fit. The estimation follows a two-step strategy similar to Gourinchas and Parker (2002).

In the first step, I estimate parameters that can be identified directly from the data or taken from the literature. These include the processes governing earnings, medical expenditures, survival probabilities, and health transitions. Estimating these components outside the model disciplines the structural estimation and avoids conflating external primitives with internally calibrated parameters.

In the second step, conditional on the first-stage estimates and fixed parameters, I estimate the remaining parameters using the Method of Simulated Moments (MSM). I select a set of empirical moments and minimize the unweighted Euclidean distance between the observed and simulated values generated by the model.

The parameters estimated in the second stage include: (i) the coefficient of relative risk aversion in the CRRA function, γ; (ii) the health-dependent preference shifters for physical and mental health, δ and ψ; (iii) the bequest motive parameters: the strength of the bequest motive, θ_{beq}, and the degree to which bequests are a luxury good, κ_{beq}; and (iv) $\underline{c}$ the government-provided consumption floor.

Together, these parameters capture the preference side of the model and are identified from the life-cycle patterns of wealth observed in the data.

5.1 First-step estimates

Exogenous Parameters: The annual discount factor β is set to 0.95 and the biannual interest rate to $r = 0.02$. These values are standard in the macroeconomic life-cycle literature and provide a benchmark consistent with prior studies.

Earnings Process: The estimation sample includes household heads aged 25–61 who report positive earnings and hourly wages above the federal minimum wage[6]. Following Storesletten et al. (2004), I estimate the earnings process by regressing the logarithm of earnings on time fixed effects, individual fixed effects, and a quadratic polynomial in age.

[6]These filters ensure that the estimation focuses on individuals who are active in the labor market, not retired yet, and whose main income source is labor earnings.

The residuals from this regression are used to construct the variance-covariance matrix of earnings shocks. I identify the parameters of the AR(1) earnings process, $\{\sigma_\kappa, \sigma_\varepsilon, \sigma_\nu, \rho\}$, by minimizing the unweighted squared distance between this empirical variance-covariance matrix and the theoretical variance-covariance matrix implied by the AR(1) structure. This procedure ensures that the persistence and variance of earnings shocks are disciplined directly by the data.

The estimated parameters of the AR(1) process are presented in Appendix C, and the regression results for equation (6) are reported in Table 1.

Table 1: Estimation results for earnings regression.

	(1)	(2)	(3)	(4)
		dependent variable: $\log(y)$		
$\mathbb{1}\{p_{t-1} = bad\}$	-0.128***		-0.122***	-0.0656***
	(0.0198)		(0.0199)	(0.0245)
$\mathbb{1}\{m_{t-1} = bad\}$		-0.0489***	-0.0415***	-0.0214
		(0.0137)	(0.0137)	(0.0146)
$\mathbb{1}\{p_{t-1} = bad \wedge m_{t-1} = bad\}$				-0.134**
				(0.0339)
Age_t	0.0753***	0.0752***	0.0756***	0.0762***
	(0.0128)	(0.0128)	(0.0128)	(0.01287)
Age_t^2	-0.000888***	-0.000895***	-0.000889***	-0.000888***
	(0.000053)	(0.000053)	(0.000053)	(0.000053)
Constant	9.112***	9.132***	9.109***	9.077***
	(0.526)	(0.526)	(0.526)	(0.525)
Observations	15,224	15,224	15,224	15,224

Household income is equivalized. Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Our main objects of interest in Table 1 are the coefficients associated with the mental and physical health dummies. When only the physical health is included—omitting mental health and the interaction term—column (1) shows that being in poor physical health is associated with an earnings reduction of 12.8 percentage points. When only mental health is included, column (2) indicates a 4.9 percentage point decline. Including both variables simultaneously, as in column (3), yields similar estimates: bad physical health is associated with a 12.2 percentage point reduction in earnings, and bad mental health a decline of 4.2 percentage points.

Introducing the interaction term between the two health dummies, however, changes the picture substantially. As shown in column (4), the coefficient on bad physical health and bad mental health fall to 6.6 and 2.1 percentage points, respectively, with the latter losing statistical significance. The interaction term, by contrast, is large and highly significant: experiencing both poor mental and physical health is associated with a 13.4 percentage points reduction in earnings. These results underscore the importance of explicitly distinguishing between mental and physical health in the earnings process—ignoring their interaction can mask the true joint impact of health on labor income.

Health Transition Process: To estimate the probability of moving across health states, I estimate a multinomial logit model. The probability of transitioning to health state j in the next period, conditional on being in health state k today, at age t, and belonging to productivity type z, is given by:

$$\Pr(h_{t+1} = j | h_t = k, z, t) = \frac{\exp\left\{\beta_j^0 + \beta_k^j \mathbb{1}\{h_t = k\} + \beta_{z,q}^j \mathbb{1}\{z = q\} + \beta_{1,t}^j t + \beta_{2,t}^j t^2\right\}}{1 + \sum_{s=1}^{5} \exp\left\{\beta_0^s + \beta_k^s \mathbb{1}\{h_t = k\} + \beta_{z,q}^s \mathbb{1}\{z = q\} + \beta_{1,t}^s t + \beta_{2,t}^s t^2\right\}} \quad (15)$$

Figure 3 presents the predicted transition probabilities for three future states: being in bad mental and physical health (top panel), good mental and physical health (middle panel), and death (bottom panel)[7]. Each panel contains four subplots corresponding to distinct initial health states. Within each subplot, curves represent each productivity type, allowing a comparison of how health dynamics differ across low-, medium-, and high-productivity individuals.

[7] Appendix D, presents results for the remaining health states.

Figure 3: Predicted health transition probabilities by initial health state

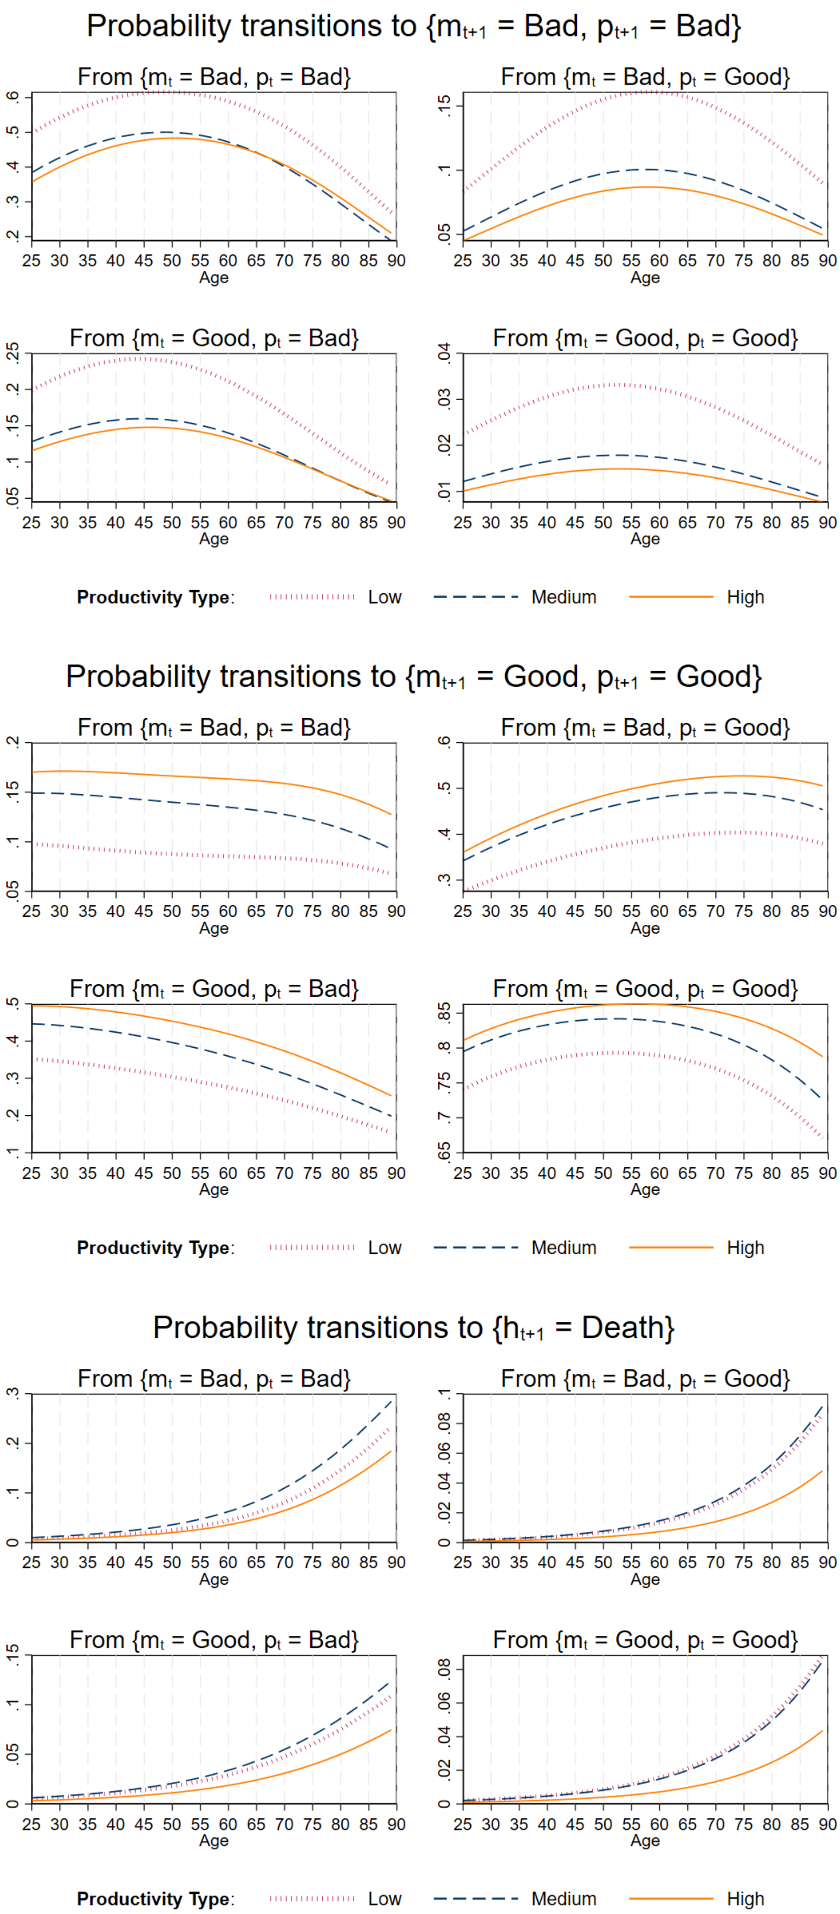

These graphs reveal a strong two-way link between health and productivity. Health status affects earnings, as shown earlier, but individual's productivity type also shapes future health transitions.

In the top panel of Figure 3, individuals with low productivity are far more likely to experience both poor mental and physical health than those with medium or high productivity, who display similar patterns. The middle panel shows a clear ranking across productivity groups in terms of the probability of recovery (for an individual who is currently ill) or staying healthy. As expected, high-productivity individuals are the most likely to remain or become healthy ($m_t = \text{good}, p_t = \text{good}$), while those with low productivity face the weakest prospect. Mortality differences are smaller. As show in the bottom panel, only the high-productivity group stands out, exhibit a noticeably lower probability of death—particularly at older ages—while medium- and low-productivity types display similar mortality risks.

A second striking feature is the persistence of health states. Individuals already in poor mental and physical health (m_t = bad, p_t = bad) are most likely to remain there, just as those currently in good health (m_t = good, p_t = good) tend to stay healthy. Transitions into mixed states are less frequent, underscoring the stability of both good and bad health states over time.

Finally, the figure highlights a systematic survival advantage for high-productivity individuals. The gap is especially pronounced when physical health is good, where their mortality risk diverges sharply from that of medium- and low-productivity groups. In contrast, medium- and low-productivity individuals exhibit nearly identical survival probabilities.

Medical expenses process: Medical expenses are modeled in two steps. First, I estimate the probability that an individual incurs positive out-of-pocket medical expenditures using a binary logit model. The likelihood of having expenses at age t, conditional on being in health state k, is given by

$$\Pr(e_t = 1 | h_t = k, t) = \frac{\exp\left\{\beta^0 + \beta_k \mathbb{1}\{h_t = k\} + \beta_{1,t} t + \beta_{2,t} t^2 + \beta_{3,k} \mathbb{1}\{h_t = k\} \times t\right\}}{1 + \exp\left\{\beta_0 + \beta^s_{1,t} t + \beta^s_{2,t} t^2 + \beta_k \mathbb{1}\{h_t = k\} + \beta_{3,k} \mathbb{1}\{h_t = k\} \times t\right\}} \tag{16}$$

Second, conditional on having positive expenditures, I model the level of spending using the following regression:

$$\begin{aligned} \ln e_{it} =& \alpha_i + \tau_t + \beta_1 t + \beta_2 t^2 + \gamma_1 \mathbb{1}\{m_t = \text{bad}\} + \gamma_2 \mathbb{1}\{p_t = \text{bad}\} + \gamma_3 \mathbb{1}\{p_t = \text{bad} \ \wedge m_t = \text{bad}\} \\ &+ \delta_1 \mathbb{1}\{p_t = \text{bad}\} \times t + \delta_2 \mathbb{1}\{m_t = \text{bad}\} \times t + \delta_3 \mathbb{1}\{m_t = \text{bad} \ \wedge p_t = \text{bad}\} \times t + \varepsilon_{it} \end{aligned} \tag{17}$$

Table 2 reports regression estimates of log out-of-pocket medical expenditures for individuals with positive spending. The results show that poor health is associated with higher medical expenditures, but the relation is stronger and more robust for mental health than for physical health. Being in bad mental health raises log expenditures by roughly 14–20 percentage points across specifications, and this estimate remains significant even after controlling for age interactions. In contrast, the relation with bad physical health is smaller and loses significance once the set of covariates is included.

Overall, the results underscore the importance of distinguishing between mental and physical health when modeling medical expenditures. Mental health problems are a more consistent and powerful driver of out-of-pocket spending, while the joint presence of both poor mental and physical health leads to especially high costs.

Figure 4 plots the predicted probability of having positive medical expenses and the corresponding spending amounts by age and health state. Panel 4a shows that individuals in poor mental and physical health spend substantially more than their healthier counterparts, although this gap narrows with age. Between ages 25 and 30, those

Table 2: Estimation results for medical expenses.

	(1)	(2)	(3)	(4)
	dependent variable: $\log(y)$			
$\mathbb{1}\{p_t = bad\}$	0.178*		0.150	0.00731
	0.0923)		(0.0929)	(0.120)
$\mathbb{1}\{m_t = bad\}$		0.196***	0.188***	0.144**
		(0.0654)	(0.0660)	(0.0717)
$\mathbb{1}\{p_t = bad \wedge m_t = bad\}$				0.304**
				(0.151)
$\mathbb{1}\{p_t = bad\} \times Age_t$	-0.002		-0.00162	0.00070
	(0.00153)		(0.00154)	(0.00197)
$\mathbb{1}\{m_t = bad\} \times Age_t$		-0.0023*	-0.00223*	-0.00155
		(0.00117)	(0.00119)	(0.00133)
$\mathbb{1}\{p_t = bad \wedge m_t = bad\} \times Age_t$				-0.00446*
				(0.0025)
Age_t	0.136**	0.139**	0.140**	0.139**
	(0.0571)	(0.0570)	(0.0571)	(0.0571)
Age_t^2	-0.000542**	-0.000554**	-0.000554**	-0.000547**
	(0.00024)	(0.00024)	(0.00024)	(0.00024)
Constant	0.916***	0.813	0.764	0.786
	(2.871)	(2.867)	(2.870)	(2.870)
Observations	40,281	40,281	40,281	40,281

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

in poor health spend about 35 percentage points more, but by ages 81–90 the difference shrinks to roughly 0.14 percentage points. Panel 4b depicts the extensive margin: mentally healthy individuals are over 10 percentage points less likely to incur medical expenses at younger ages (25-40), yet these differences also diminish with age, converging to similar probabilities late in life.

Taken together, these results show that the relationship between health and medical spending is strong but highly age-dependent. On the intensive margin, individuals in poor mental and physical health spend about 35 percentage points more than healthy individuals in early adulthood (ages 25-30), yet this gap monotonically narrows with age, falling to nearly zero by age 90. On the extensive margin, mentally healthy individuals are over 10 percentage points less likely to incur medical expenses at young ages (25-40), but this difference also disappears later in life as out-of-pocket spending becomes nearly universal. Overall, poor mental health emerges as a key driver of

Figure 4: Predicted annual medical expenses by age and health state

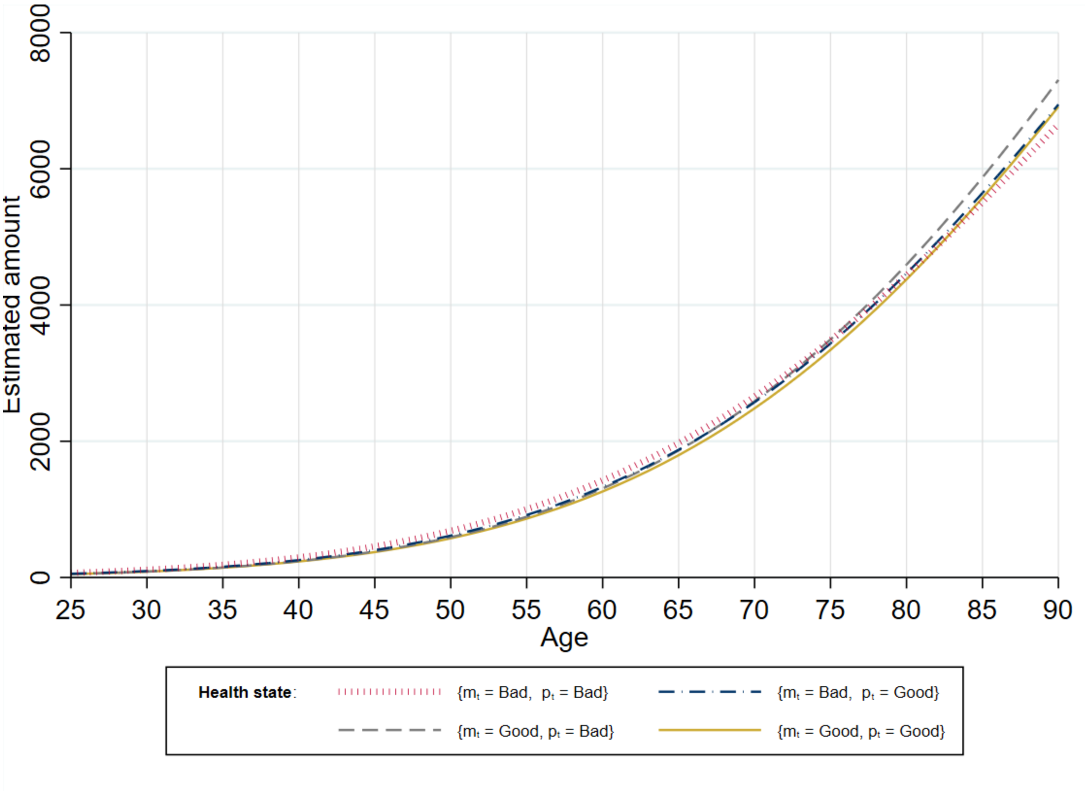

(a) Estimated amount of annual medical expenses by health state

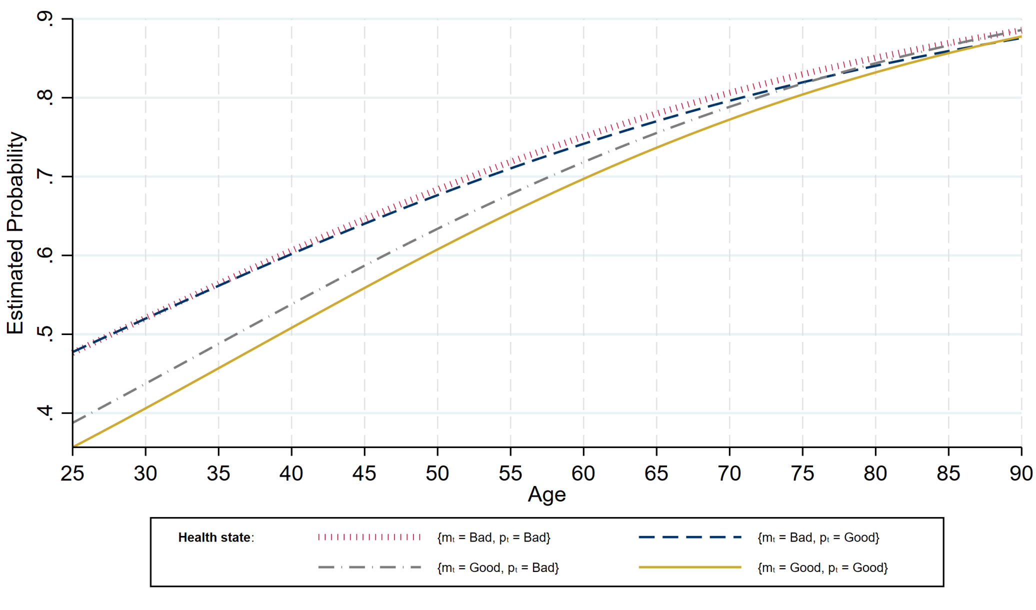

(b) Estimated probability of having positive medical expenses

higher medical costs at younger ages, while the influence of health status on both the level and likelihood of spending fades with age.

5.2 Second-step estimates

In the second step, I estimate the risk-aversion parameter, the health-dependent marginal utility shifters, the consumption floor, and the bequest parameters so that the model matches key moments of the wealth distribution. The estimation targets average wealth by age and health status, using ten six-year age bins[8] and four health categories defined by the combination of good and bad states across the two health dimensions.

The estimation proceeds in four steps. First, given an initial guess for the parameters, I solve the finite-horizon life-cycle model to obtain the optimal consumption and saving policies. Second, I simulate individuals over the life cycle using these policy functions. Third, I compute the average wealth in each age-health bin and construct the objective function as the sum of squared differences between the simulated and empirical averages. Finally, I use the Nelder-Mead algorithm to search over the parameter space for the set of values that minimizes this distance.

5.3 Identification

Targeting wealth profiles by age and health status is consistent with both data availability and the model's theoretical structure. Wealth accumulation results from life-cycle consumption and saving choices, which are in turn jointly determined by preferences, health risks, and bequest motives. Conditioning on both age and health allows the estimation to capture heterogeneity in savings behavior driven by differences in health shocks and health-dependent preferences.

In a non-linear model like this, all parameters potentially affect all moments, but some moments are more informative for particular parameters. In this section, I outline the main sources of identification for each parameter from the model.

[8] All groups span six years except the three post-retirement groups, which each aggregates eight ages.

Risk aversion parameter, γ: The curvature of the CRRA utility function governs how sensitive consumption is to risk and intertemporal substitution. A higher γ implies stronger precautionary saving in response to uncertainty and smoother consumption paths across time and health states. In estimation, γ is primarily disciplined by the overall shape of wealth accumulation over the life cycle—that is, by the extent to which households build precautionary buffers against income, health, and mortality risks.

Health-dependent preference shifters, δ, and ψ: These parameters govern how the marginal utility of consumption varies with health status. Their identification exploits differences in wealth profiles across health groups that go beyond earnings and mortality differences. In the absence of these shifters, individuals would equate marginal utility across health states and save accordingly. The systematic wealth gaps observed across health conditions suggest that utility shifts with health, allowing the estimation to recover δ and ψ from these differences.

Bequest parameters, θ_{beq}, and κ_{beq}: The strength of the bequest motive, θ_{beq}, and the curvature parameter, κ_{beq}, are identified from late-life wealth profiles. The persistence of substantial assets at advanced ages—when survival probabilities are low—cannot be rationalized solely as precautionary saving. A positive θ_{beq} accounts for continued wealth holding despite high mortality risk, while κ_{beq} determines whether bequests behave as a luxury good (relevant only for the wealthy) or as a broader phenomenon across the population. Hence, the shape of the wealth distribution late in life provides key information for identifying both parameters.

Consumption floor, $\underline{c}$: The minimum consumption level financed by the government represents the insurance available to the poorest individuals. Its value is identified from the lower tail of the consumption and wealth distributions. A higher floor reduces precautionary saving among low-wealth individuals and compresses inequality in consumption responses. In the data, the limited asset accumulation and higher consumption sensitivity of poor individuals to transitory income shocks are the key moments that discipline this parameter.

In summary, identification in this framework relies on distinct features of the wealth distribution by age and health status. The risk aversion parameter γ is disciplined by overall wealth accumulation over the life cycle; the health-dependent utility shifters δ and ψ by wealth differences across health states; the bequest parameters θ_{beq} and κ_{beq} by the persistence of wealth into old age; and the consumption floor $\underline{c}$ by the behavior of households at the bottom of the distribution. Together, these parameters enable the model to match both average life-cycle profiles and the heterogeneity in saving behavior observed in the data.

5.4 Estimation results and model fit

In my estimates, I obtain a coefficient of relative risk aversion that is considerably lower than the values typically found and assumed in the literature. Specifically, the estimated parameter $\gamma = 0.319$, whereas most studies in macroeconomics estimate or assume values between 2 and 4. The unusually low value in this case may reflect specific features of the data, particularly the difference in the curvature across the different health states.

Table 3: Estimated structural parameters

Parameter	Definition	Estimates
γ	Coefficient of relative risk aversion	0.319
δ	Utility shifter for physical health	1.335
ψ	Utility shifter for mental health	0.867
θ_{beq}	Bequest weight	5.978
κ_{beq}	Curvature of the bequest	374,736
$\underline{c}$	Consumption floor	11,427

Regarding the health-dependent preferences parameters, the estimates yield $\psi = 0.867$ for mental health and $\delta = 1.335$ for physical health. Both imply that marginal utility increases when individuals are in poor health. This finding is consistent with

Kools and Knoef (2019) and aligns with a broader literature documenting a positive relationship between poor health and the marginal utility of consumption (Lillard and Weiss (1997), Rust and Phelan (1997), Edwards (2008), Ameriks et al. (2020)).

For bequest motives, the estimated parameter $\theta_{beq} = 5.978$ indicates that individuals derive utility from leaving bequests, while the large value of $\kappa_{beq} = 374,736$ suggests that bequests operate primarily as a luxury good—relevant mainly for wealthier households. The estimated government consumption floor corresponds to a biannual transfer of $\underline{c} = 11,427$ dollars.

Figure 5 compares the life-cycle wealth profiles from the PSID with those generated by the estimated model. The model reproduces the targeted patterns closely, capturing both the average wealth accumulation and the differences across health states. In Appendix E, I offer more details on the model's fit for aggregate wealth moments.

Taken together, the estimation results suggest a preference structure characterized by relatively low curvature in consumption but strong sensitivity of marginal utility to health conditions and bequest motives. The low estimate of γ implies limited overall risk aversion once health-dependent shifters are incorporated, while the positive values of δ and ψ confirm that poor health raises the marginal value of consumption, consistent with previous empirical findings.

The bequest parameters further indicate a meaningful, though concentrated, bequest motive that operates mainly among wealthier households, as reflected by the large luxury component. Overall, the estimates support a framework in which health status and intergenerational motives jointly shape saving and consumption behavior over the life cycle.

Figure 5: Targeted Moments: Comparison between model and data by health state

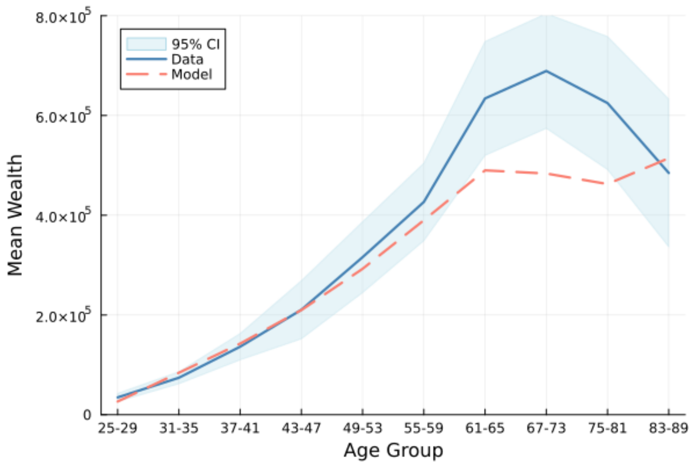

(a) m_t = Good, p_t = Good

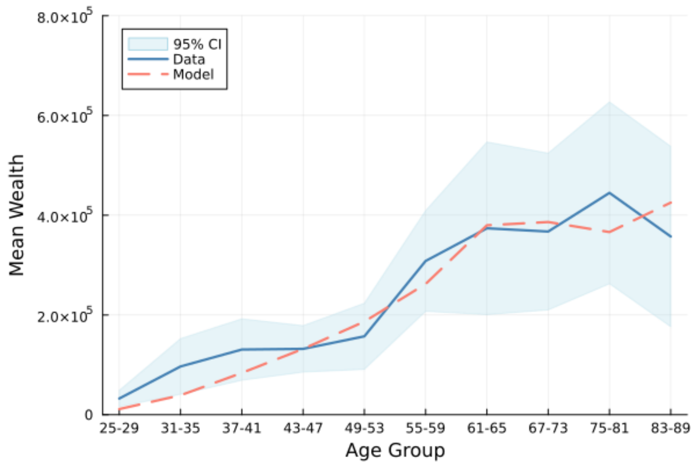

(b) m_t = Bad, p_t = Good

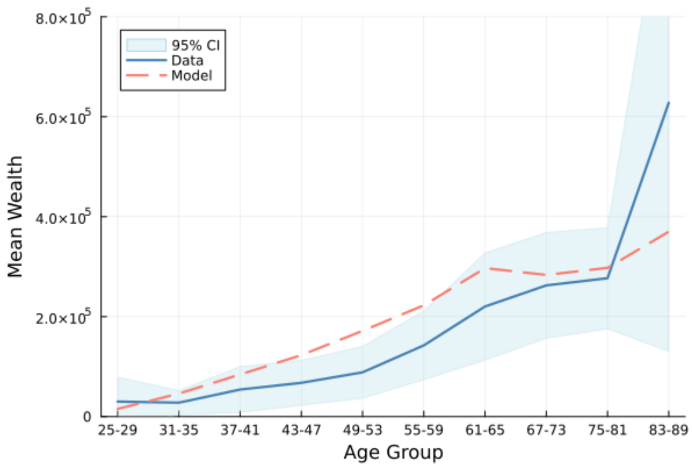

(c) m_t = Good, p_t = Bad

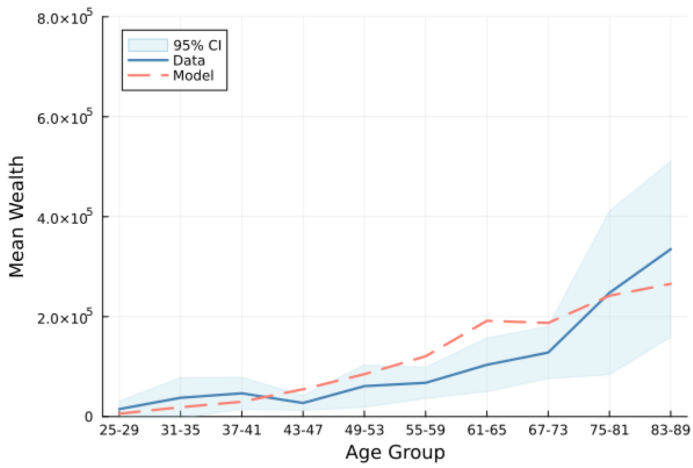

(d) m_t = Bad, p_t = Bad

6 Health-Dependent Preferences, Consumption and Savings

In this section, I use the estimated model to quantify how health-dependent preferences shape life-cycle consumption and savings. The analysis builds on Russo (2023), who studies the aggregate implications of health-dependent marginal utility, but extends that framework by explicitly distinguishing between mental and physical health. This distinction allows me to assess the relative importance of each health dimension in driving consumption-saving decisions and to evaluate how the marginal utility of consumption responds to poor mental health, poor physical health, or both simultaneously.

Using the baseline calibration, I re-solve the model under a series of counterfactual scenarios that alter the link between health and preferences. Specifically, I consider three cases: (i) physical-health dependence removed ($\delta = 0$); (ii) mental-health dependence removed ($\psi = 0$); and (iii) both shifters set to zero ($\delta = \psi = 0$), thereby fully removing health effects on preferences.

For each counterfactual, I simulate the life cycle of 50,000 individuals and compare the resulting consumption and savings profiles to those in the baseline model. This exercise isolates the separate and joint contributions of each health dimension to life-cycle economic behavior.

In the baseline scenario, the preference parameters take values of $\psi = 0.87$ for mental health and $\delta = 1.33$ for physical health. These positive coefficients generate two opposing forces in saving decisions. On one hand, poor health raises the marginal utility of current consumption, encouraging higher spending and reducing asset accumulation; on the other, the anticipation of future health deterioration increases precautionary motives, inducing higher savings at earlier ages.

Figure 6: Health dependent preferences and optimal wealth

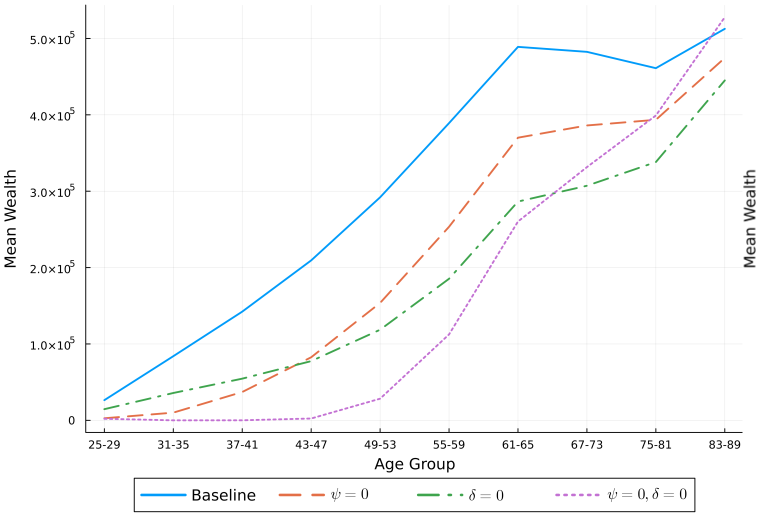

(a) m_t = Good, p_t = Good

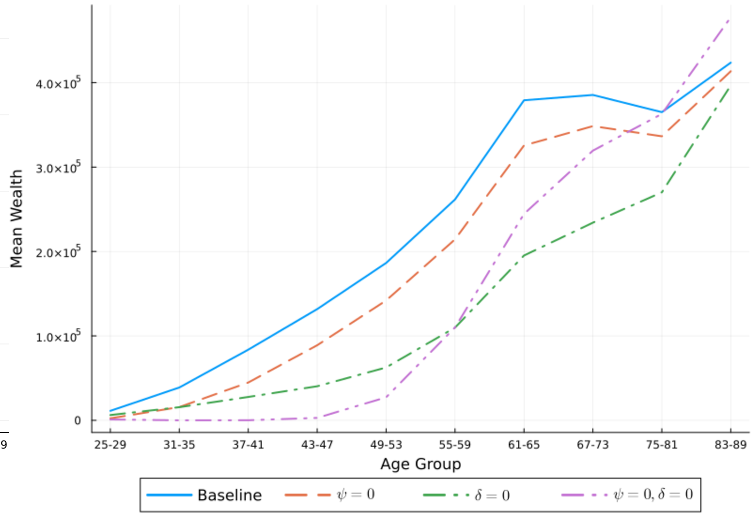

(b) m_t = Bad, p_t = Good

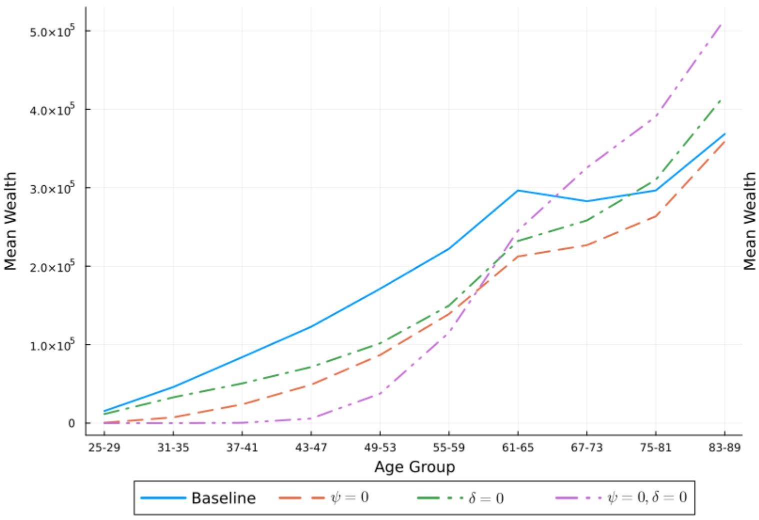

(c) m_t = Good, p_t = Bad

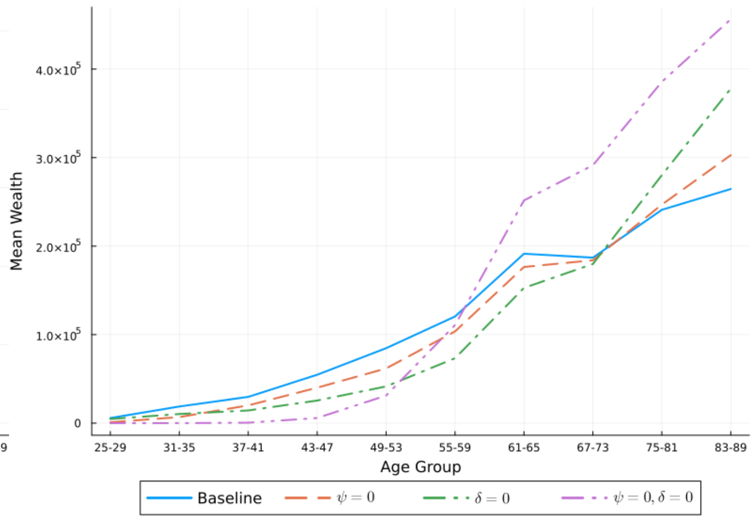

(d) m_t = Bad, p_t = Bad

Figure 6 plots the life-cycle profiles of average wealth across age bins, with each panel corresponding to a specific health state. The solid line represents the baseline scenario, while the dashed lines show the counterfactuals in which health-dependent shifters are partially or fully removed.

At younger ages, removing the marginal-utility shifters lowers the incentive to save, as individuals no longer anticipate higher utility from consumption in future periods of poor health. Consequently, savings decline relative to the baseline, especially among those in good physical health.

Later in life, the pattern reverses. Among individuals in poor health, removing the shifters eliminates the strong preference for immediate consumption, leading to higher savings and greater end-of-life wealth in the counterfactuals. The effect is most pronounced for individuals who experience poor health in both dimensions.

Taken together, these results show that health-dependent preferences are a key driver of intertemporal allocation. They alter the timing of saving and consumption by increasing the marginal value of resources in bad health. Removing this channel weakens precautionary savings among the healthy but raises them among those already in poor health. Overall, the economic role of health-dependent preferences is strongly state-dependent—amplifying differences in behavior across health conditions and over the life cycle.

7 Counterfactuals

7.1 Worsening initial distribution of mental health

A growing concern in public health is that younger cohorts are reaching adulthood with worse mental health than previous generations. Recent studies in economics confirm this trend (Blanchflower et al. (2024b); Blanchflower et al. (2024a); Blanchflower and Bryson (2024)). To assess its macroeconomic implications, I simulate a counterfactual scenario in which individuals begin life with a worse initial distribution of mental health than in the baseline.

The motivation for this exercise stems from the persistence of mental health states and the structure of preferences. Since $\psi > 0$, poor mental health raises the marginal utility of current consumption, altering the trade-off between immediate consumption and precautionary saving. Deteriorating initial conditions therefore shift consumption–saving behavior early in life, with downstream consequences for wealth accumulation and bequest patterns.

To isolate the role of initial conditions, all other elements of the model—transition probabilities, prices, and parameters—are held constant. I reassign 20% of individuals who would otherwise begin in good mental health at age 25 (m_t = Good) to bad mental health instead. This reallocation is random and independent of physical health status p_t, so that the deterioration affects both physically healthy and unhealthy individuals proportionally. Table 4 reports the resulting health distribution at age 25 for each productivity type and compares it with the baseline.

The counterfactual increases the share of individuals in poor mental health at age 25

Table 4: Health distribution at age 25 by productivity type.

	Health state	Type 1	Type 2	Type 3
Baseline	m_t = Bad, p_t = Bad	0.04	0.04	0.06
	m_t = Bad, p_t = Good	0.36	0.29	0.31
	m_t = Good, p_t = Bad	0.04	0.03	0.01
	m_t = Good, p_t = Good	0.56	0.64	0.62
Counterfactual	m_t = Bad, p_t = Bad	0.05	0.04	0.06
	m_t = Bad, p_t = Good	0.47	0.42	0.44
	m_t = Good, p_t = Bad	0.03	0.02	0.01
	m_t = Good, p_t = Good	0.45	0.51	0.50

by 13 percentage points. As shown in Figure 7, this gap relative to the baseline narrows quickly and virtually disappears within roughly eight years, after which the two distributions converge. For physical health, Panel 7b shows that the larger fraction of individuals with poor mental health slightly raises the share of those in bad physical health—by about 0.4 percentage points—but the effect is small and short-lived, leaving no lasting differences between the two scenarios.

Figure 7: Share of ill individuals by age

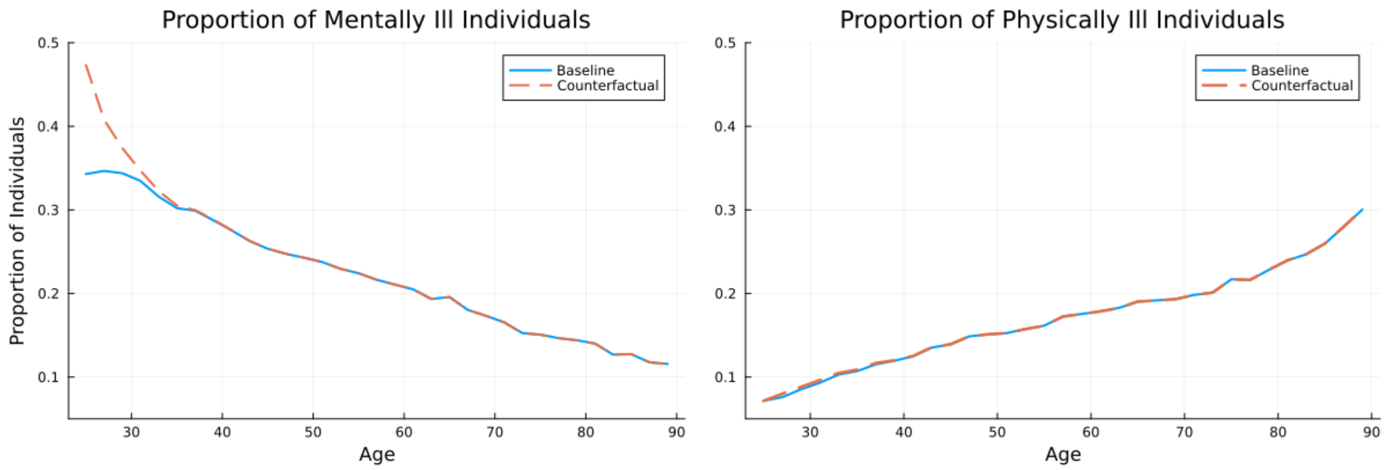

(a) Mental health

(b) Physical health

On the earnings side, the effect of worsening the initial distribution of mental health is minimal. In the model, bad mental health alone produces only a modest earnings loss; substantial penalties arise primarily when mental and physical health problems coincide. At age 25, the prevalence of poor physical health is relatively low. Consequently, the counterfactual reallocation moves only a small fraction of individuals into the joint state. As a result, the earnings profiles in Figure 8a, under the counterfactual, are virtually identical to those in the baseline.

Figure 8: Aggregate Earnings, Consumption, and Wealth by age.

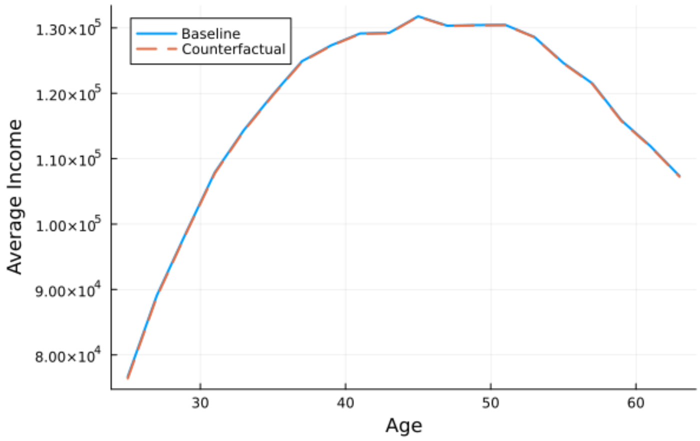

(a) Earnings

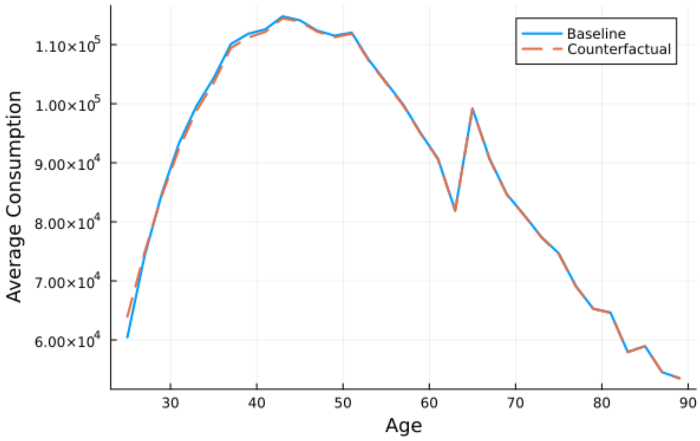

(b) Consumption

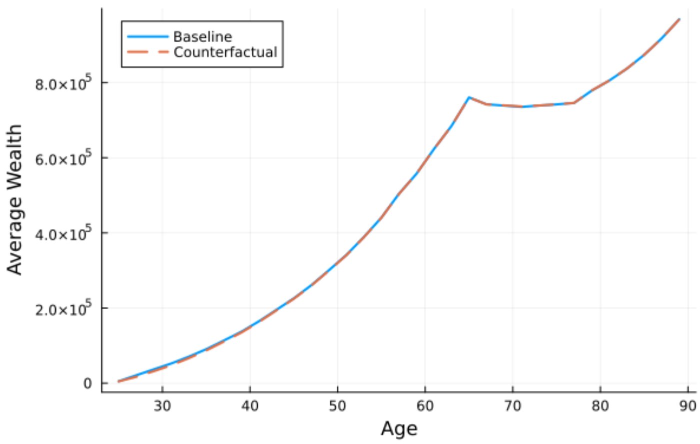

(c) Wealth

While earnings remain virtually unchanged, health-dependent preferences alter the timing of consumption over the life cycle. As shown in Figure 8b, worsening the initial mental health distribution raises the share of young entrants with higher marginal utility of current consumption (since $\psi > 0$ when $m_t = \text{Bad}$). This higher return to present spending leads agents to front-load consumption, reducing early-life savings. At age 25, consumption in the counterfactual is 5.8 percentage points higher than in the baseline, but lower thereafter.

Two mechanisms drive this pattern: (i) it stems from differences in preferences rather than earnings, since earnings are nearly identical across scenarios, the Euler condition implies that higher marginal utility today weakens the incentive to shift resources forward, compressing midlife consumption; (ii) even modest changes in early-life consumption have large implications for wealth accumulation. Between ages 25 and 31, asset holdings fall by more than 12 percentage points in the counterfactual, signaling a sharp reduction in early wealth building. Additional details on the joint evolution of consumption and wealth by health state are provided in Appendix F.

Overall, the experiment shows that worsening initial mental health has negligible effects on earnings but substantially reshapes the timing of consumption and saving. Because mental health affects preferences rather than income, the main adjustment occurs through intertemporal substitution: agents consume more early in life and save less for the future. This highlights how initial conditions can significantly shift the allocation of resources across the life cycle—not by altering economic capacity, but by adjusting the balance between present and future consumption.

7.2 Worsening initial distribution of physical health

Epidemiological evidence documents a rise in chronic conditions among younger cohorts in the United States—such as obesity (Temple, 2022) and multiple comorbidities (Watson et al., 2025). This deterioration reflects not only higher disease incidence but also behavioral changes, including declining physical activity levels citep SFGSCRBS2024. Motivated by these trends, I construct a counterfactual experiment that worsens the initial distribution of physical health at the onset of adulthood.

As in the previous exercise, all other components of the model remain fixed. I modify only the initial physical health distribution at age 25, reallocating 20% of individuals who would otherwise be in good physical health ($p_t =$ Good) into the bad physical health. This reallocation is independent of mental health status. Table 5 reports the resulting distribution of health at age 25 for each productivity group.

Table 5: Health distribution at age 25 by productivity type.

	Health state	Type 1	Type 2	Type 3
Baseline	$m_t = \text{Bad}, p_t = \text{Bad}$	0.04	0.04	0.06
	$m_t = \text{Bad}, p_t = \text{Good}$	0.36	0.29	0.31
	$m_t = \text{Good}, p_t = \text{Bad}$	0.04	0.03	0.01
	$m_t = \text{Good}, p_t = \text{Good}$	0.56	0.64	0.62
Counterfactual	$m_t = \text{Bad}, p_t = \text{Bad}$	0.11	0.10	0.12
	$m_t = \text{Bad}, p_t = \text{Good}$	0.29	0.23	0.25
	$m_t = \text{Good}, p_t = \text{Bad}$	0.15	0.16	0.13
	$m_t = \text{Good}, p_t = \text{Good}$	0.45	0.51	0.50

This experiment increases the share of individuals in poor physical health at age 25 by 18 percentage points.[9]. Figure 9 shows the dynamics of mental and physical health over the life-cycle. Although the intervention generates a larger initial rise in physical illness, the persistence of the shock is similar to that of the mental health counterfactual: the gap relative to the baseline vanishes after roughly eight years.

A key difference, however, is that the deterioration in physical health spills over into mental health in subsequent periods. The share of mentally ill individuals rises by 1.8, 1.7, and 1.0 percentage points at ages 27, 29, and 31, respectively. These results suggest that, early in the life cycle, poor physical health exerts a stronger influence on mental health than the reverse.

On the earnings side, the counterfactual produces an immediate 1.8 percentage points decline at age 25, which quickly lowers to 0.8 percentage points at age 27 and 0.3 percentage points at age 29. The stronger initial effect relative to the mental health experiment reflects the lower baseline prevalence of poor physical health (7% versus 34% for poor mental health). As a result, worsening the physical health distribution substantially increases the likelihood of individuals experiencing both poor physical and mental health, amplifying early-life earnings losses.

[9]This number is larger than in the previous experiment because, at age 25, the proportion of individuals in good physical health exceeds that of good mental health.

Figure 9: Share of ill individuals by age

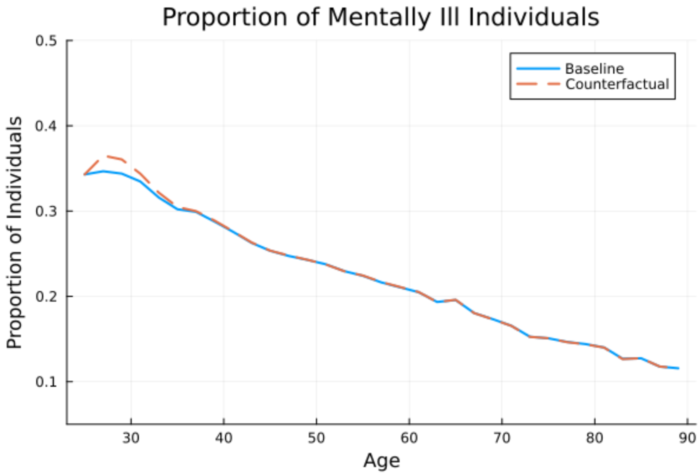

(a) Mental health

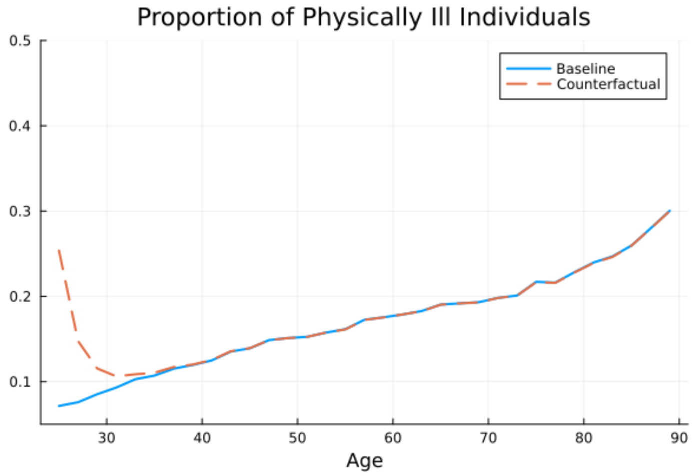

(b) Physical health

Because the reallocation moves previously healthy individuals into a state associated with larger income penalties—6.5 percentage points if they were already mentally ill and 13.4 percentage points if mentally healthy—the impact on consumption is comparatively muted. On one hand, individuals have a higher incentive to consume given that $\delta > 0$; on the other, lower earnings constrain available resources. As shown in Figure 10b, the net effect is modest: consumption rises by about 3 percentage points initially but declines by roughly 0.2 percentage points in subsequent periods.

The result is a sizable decline in wealth accumulation. At age 27, the aggregate wealth is almost 20 percentage points lower than in the baseline, and asset holdings drop by about 10 percentage points between ages 25 and 31.

In summary, worsening the initial distribution of physical health produces the same substitution pattern between current consumption and savings as in the previous experiment. However, since the income effect is stronger this time, the rise in early-life consumption effect is more modest, while the reduction in wealth remains equally significant.

Figure 10: Aggregate Earnings, Consumption, and Wealth by age.

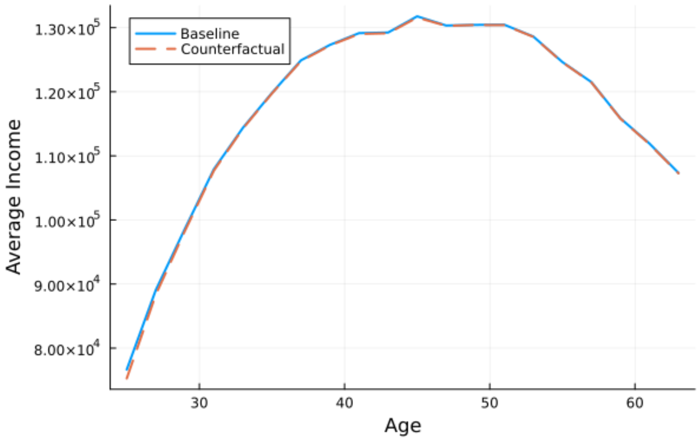

(a) Earnings

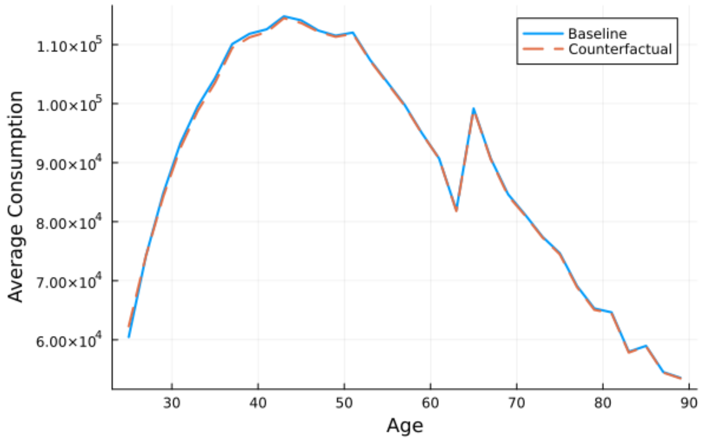

(b) Consumption

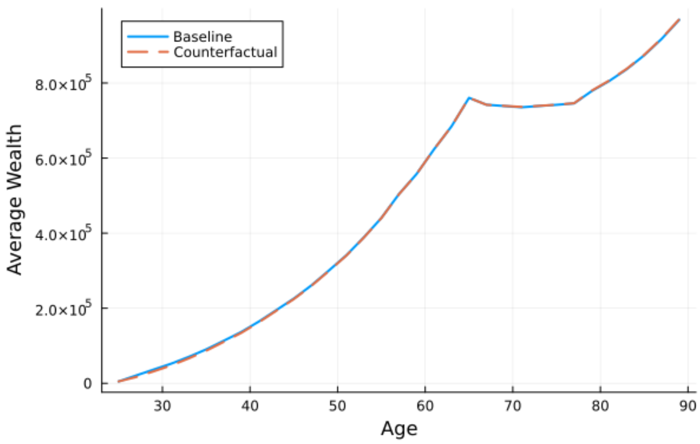

(c) Wealth

8 The Cost of Bad Mental and Physical Health

In this section, I quantify the monetary losses associated with bad mental and physical health. For each individual, I compare observed outcomes with a counterfactual scenario in which the same individual experiences no adverse health shocks—mental, physical, or both—while facing the same environment. This comparison isolates the pecuniary cost of poor health over the life cycle. I also decompose these losses by productivity type of assess how the pecuniary burden of bad health varies across individuals with different earnings potential.

Table 6: Biannual Monetary costs due to bad health

	All	θ_L	θ_M	θ_H
	No mental health shocks			
Over the life-cycle (25-death)				
Share of time in bad mental health	24%	30%	24%	21%
Income losses + total medical costs	$1,163	$ 863	$ 1,058	$ 1,664
(Percentage of group's average earnings)	(1.1%)	(1.7%)	(1.2%)	(1.0%)
Over the working age (25-63)				
Share of time in bad mental health	27%	33 %	26 %	24 %
Income losses + total medical costs	$ 1,356	$ 928	$ 1,199	$ 2,086
(Percentage of group's average earnings)	(1.3%)	(1.9%)	(1.3%)	(1.3%)
	No physical health shocks			
Over the life-cycle (25-death)				
Share of time in bad physical health	16%	21%	15%	12%
Income losses + total medical costs	$ 1,336	$ 998	$ 1,239	$ 1,859
(Percentage of group's average earnings)	(1.3%)	(2.0%)	(1.4%)	(1.1%)
Over the working age (25-63)				
Share of time in bad physical health	13%	18 %	13 %	10 %
Income losses + total medical costs	$ 1,528	$ 1,067	$ 1,384	$ 2,266
(Percentage of group's average earnings)	(1.5%)	(2.2%)	(1.5%)	(1.4%)
	No mental or physical health shocks			
Over the life-cycle (25-death)				
Share of time in bad mental or physical health	33%	40%	32 %	28 %
Income losses + total medical costs	$ 3,208	$ 2,608	$ 3,021	$ 4,162
(Percentage of group's average earnings)	(3.2%)	(5.3%)	(3.4%)	(2.5%)
Over the working age (25-63)				
Share of time in bad mental or physical health	33%	39 %	32 %	29 %
Income losses + total medical costs	$ 2,616	$ 1,908	$ 2,406	$ 3,722
(Percentage of group's average earnings)	(2.6%)	(3.9%)	(2.7$)	(2.3%)

Average of biannual earnings for whole sample is $100,494 and $49,343, $89,855, and $164,850 for the low, medium, and high productivity subgroups, respectively.

Following De Nardi et al. (2024), I compute the monetary costs of poor health by comparing income net of medical expenditures in the baseline and counterfactual scenarios. Let y_i^0 and y_i^C denote individual i's net income at time t in the baseline and counterfactual cases, respectively. The difference between the two represents the monetary costs of bad health in period t. Lifetime costs are obtained as the present value of these differences over the life cycle, using and biannual discount rate of $r = 0.02$.

Table 6 summarizes the results from the three counterfactual experiments that quantify the monetary losses from mental and physical health shocks across productivity types. Each panel corresponds to a scenario in which one or both sources of health risk are removed. The top panel eliminates mental health shocks, isolating the role of physical health shocks in driving income losses and medical expenditures. The middle panel removes physical health shocks while keeping mental health shocks active. The bottom panel suppresses both sources of risk, providing a benchmark for the combined effect of health shocks.

For each counterfactual, I calculate biannual monetary losses over two horizons—(i) the entire life cycle (ages 25 to death) and (ii) the working stage (ages 25–63). For each productivity group—low (θ_L), medium (θ_M), and high (θ_H)—the table reports the share of time spent in bad health, total income and medical expenditure losses, and these losses as a percentage of average earnings. This structure enables a direct comparison of the economic burden of mental and physical health shocks across productivity types and life cycle stages.

In the baseline, individuals spend about 24 percent of their lives in poor mental health, generating combined income and medical expenditure losses of roughly $1,163—or 1.1 percent of average earnings. Similarly, physical health problems account for about 16 percent of life years in poor health in the baseline, with total monetary losses of around $1,336, 1.3 percent of average earnings. Thus, both dimensions of health impose comparable aggregate costs, even though physical health problems are less frequent.

When both mental and physical health shocks are jointly removed, individuals remain healthy throughout the entire life cycle. In the baseline, they would have spent about 33% of heir lifetime in either bad mental or bad physical health. The removal of both types of shocks eliminate losses equivalent to roughly $3,208, or 3.2 percent of average earnings. This result reveals the strong complementarity between mental and physical health risks: their joint presence amplifies total economic costs well beyond the sum of their individual effects.

Finally, the burden of poor health varies systematically with productivity. Low-productivity individuals face the highest relative costs, losing up to 5.3 percent of average earnings in the baseline, compared with about 2.5 percent for high-productivity individuals. This gradient suggests that avoiding health shocks yields proportionally greater economic gains for lower-productivity workers, for whom financial consequences of poor health are more severe.

9 Conclusion

This paper develops a structural life-cycle model that explicitly distinguishes between mental and physical health, incorporates medical expenses, and allows for health-dependent preferences and productivity heterogeneity. Using data from the PSID and MEPS, I document sharp differences in the dynamics of mental and physical health across the life cycle, as well as in the age profiles for medical spending. Mental health problems are concentrated at younger ages and improve through midlife, whereas physical health deteriorates steadily with age. Medical expenditures rise with age, and younger individuals in poor health face both higher probabilities of incurring costs and larger out-of-pocket payments. These facts highlight the need to treat mental and physical health as distinct sources of economic risk.

Earnings regressions indicate that poor mental and physical health each reduce earnings, but their joint occurrence results in a particularly large and statistically significant penalty. Moreover, productivity and health interact strongly: individuals with higher productivity not only earn more but also enjoy more favorable health transitions and lower mortality. Introducing productivity types into the model disciplines these patterns and links economic opportunity to health outcomes.

The estimated structural parameters indicate that the marginal utility of consumption increases in periods of poor health, both mental and physical, consistent with health-dependent preferences. Bequest parameters indicate that individuals retain wealth late in life not only for precautionary reasons but also due to a non-negligible, luxury-driven bequest motive. Medical expenses further amplify wealth decumulation at older ages, interacting with preferences and bequest motives to shape late-life saving behavior.

I then quantify the monetary burden of poor health by comparing each individual to a counterfactual version of themselves with no adverse health shocks. Mental and physical health shocks impose comparable economic costs when considered separately—about 1 to 1.3 percent of lifetime earnings—but their joint presence amplifies total losses to roughly 3.2 percent. These costs are relatively higher for low-productivity individuals, who lose over twice as much as their high-productivity counterparts, underscoring the unequal pecuniary burden of poor health.

Counterfactual experiments that worsen initial health at age 25 reveal that early-life health deficit primarily shifts the timing of consumption rather than long-run income. Individuals entering adulthood in poor mental or physical health increase early consumption and reduce savings, creating temporary wealth gaps that narrow later in life.

Taken together, the results show that distinguishing between mental and physical health adds substantial richness to the analysis of health and economic behavior. Treating health as multidimensional allows the model to capture heterogeneity in key outcomes—earnings, savings, and welfare—that would otherwise be obscured by a one-dimensional health index. Accounting for these two components also clarifies how health risks interact with productivity and preferences to shape life-cycle inequality. Future work could extend this framework to examine heterogeneity in access to healthcare, the distinction between preventive and curative spending, and the role of insurance and policy design in mitigating the long-term costs of poor health.

References

Abramson, B., Boerma, J., and Tsyvinski, A. (2025). Macreoconomics of mental health.

Ameriks, J., Briggs, J., Caplin, A., Shapiro, M. D., and Tonetti, C. (2020). Long-term-care utility and late-in-life saving. *Journal of Political Economy*, 128(6):2375–2451.

Ascarza-Mendoza, D., Ruelas-Huanca, W., and Velasquez, C. (2024). On the dynamics of mental health.

Blanchflower, D. G. and Bryson, A. (2024). The mental health of the young in latin america. Working Paper 33111, National Bureau of Economic Research.

Blanchflower, D. G., Bryson, A., and Bell, D. N. (2024a). The declining mental health of the young in the uk. Working Paper 32879, National Bureau of Economic Research.

Blanchflower, D. G., Bryson, A., Lepinteur, A., and Piper, A. (2024b). Further evidence on the global decline in the mental health of the young. Working Paper 32500, National Bureau of Economic Research.

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

Capatina, E. and Keane, M. (2025). Health shocks, health insurance, human capital, and the dynamics of earnings and health.

Capatina, E., Keane, M., and Maruyama, S. (2020). Health shocks and the evolution of earnings over the life-cycle.

Cornelius, B. L. R., Groothoff, J. W., van der Klink, J. J. L., and Brouwer, S. (2013). The performance of the K10, K6 and GHQ-12 to screen for present state DSM-IV disorders among disability claimants. *BMC Public Health*, 13:128.

De Nardi, M. (2004). Wealth inequality and intergenerational links. *The Review of Economic Studies*, 71(3):743–768.

De Nardi, M., Fella, G., and Paz-Pardo, G. (2019). Nonlinear household earnings dynamics, self-insurance, and welfare. *Journal of the European Economic Association*, 18(2):890–926.

De Nardi, M., Pashchenko, S., and Porapakkarm, P. (2024). The lifetime costs of bad health. *The Review of Economic Studies*, page rdae080.

De Nardi, M., French, E., and Jones, J. (2010). Why do the elderly save? the role of medical expenses. *Journal of Political Economy*, 118(1):39–75.

Edwards, R. D. (2008). Health risk and portfolio choice. *Journal of Business & Economic Statistics*, 26(4):472–485.

French, E. (2005). The effects of health, wealth, and wages on labour supply and retirement behaviour. *The Review of Economic Studies*, 72(2):395–427.

Furukawa, T. A., Kessler, R. C., Slade, T., and Andrews, G. (2003). The performance of the k6 and k10 screening scales for psychological distress in the australian national survey of mental health and well-being. *Psychological Medicine*, 33(2):357–362.

Gourinchas, P.-O. and Parker, J. A. (2002). Consumption over the life cycle. *Econometrica*, 70(1):47–89.

Hosseini, R., Kopecky, K., and Zhao, K. (2025). How important is health inequality for lifetime earnings inequality.

Hosseini, R., Kopecky, K. A., and Zhao, K. (2022). The evolution of health over the life cycle. *Review of Economic Dynamics*, 45:237–263.

Jolivet, G. and Postel-Vinay, F. (2024). A structural analysis of mental health and labour market trajectories. *The Review of Economic Studies*, page rdae071.

Kessler, R. C., Andrews, G., Colpe, L. J., Hiripi, E., Mroczek, D. K., Normand, S. L. T., Walters, E. E., and Zaslavsky, A. M. (2002). Short screening scales to monitor population prevalences and trends in non-specific psychological distress. *Psychological Medicine*, 32(6):959–976.

Kessler, R. C., Barker, P. R., Colpe, L. J., Epstein, J. F., Gfroerer, J. C., Hiripi, E., Howes, M. J., Normand, S.-L. T., Manderscheid, R. W., Walters, E. E., and Zaslavsky, A. M. (2003). Screening for serious mental illness in the general population. *Archives of General Psychiatry*, 60(2):184–189.

Kessler, R. C., Green, J. G., Gruber, M. J., Sampson, N. A., Bromet, E., Cuitan, M., Furukawa, T. A., Gureje, O., Hinkov, H., Hu, C.-Y., Lara, C., Lee, S., Mneimneh, Z., Myer, L., Oakley-Browne, M., Posada-Villa, J., Sagar, R., Viana, M. C., and Zaslavsky, A. M. (2010). Screening for serious mental illness in the general population with the k6 screening scale: results from the who world mental health (wmh) survey initiative. *International Journal of Methods in Psychiatric Research*, 19(Suppl 1):4–22.

Kools, L. and Knoef, M. (2019). Health and consumption preferences; estimating the health state dependence of utility using equivalence scales. *European Economic Review*, 113:46–62.

Lillard, L. A. and Weiss, Y. (1997). Uncertain health and survival: Effects on end-of-life consumption. *Journal of Business & Economic Statistics*, 15(2):254–268.

Low, H. and Pistaferri, L. (2015). Disability insurance and the dynamics of the incentive insurance trade-off. *American Economic Review*, 105(10):2986–3029.

Mahler, L. and Yum, M. (2024). Lifestyle behaviors and wealth-health gaps in germany. *Econometrica*, 92(5):1697–1733.

Palumbo, M. G. (1999). Uncertain medical expenses and precautionary saving near the end of the life cycle. *The Review of Economic Studies*, 66(2):395–421.

Prochaska, J. J., Sung, H.-Y., Max, W., Shi, Y., and Ong, M. (2012). Validity study of the k6 scale as a measure of moderate mental distress based on mental health treatment need and utilization. *International Journal of Methods in Psychiatric Research*, 21(2):88–97.

Rouwenhorst, K. G. (1995). *10 Asset Pricing Implications of Equilibrium Business Cycle Models*, pages 294–330. Princeton University Press, Princeton.

Russo, N. (2023). Health-dependent preferences, consumption, and insurance.

Rust, J. and Phelan, C. (1997). How social security and medicare affect retirement behavior in a world of incomplete markets. *Econometrica*, 65(4):781–831.

Storesletten, K., Telmer, C. I., and Yaron, A. (2004). Consumption and risk sharing over the life cycle. *Journal of Monetary Economics*, 51(3):609–633.

Temple, N. J. (2022). The origins of the obesity epidemic in the usa–lessons for today. *Nutrients*, 14(20).

Umucu, E., Fortuna, K., Jung, H., Bialunska, A., Lee, B., Mangadu, T., Storm, M., Ergun, G., Mozer, D. A., and Brooks, J. (2022). A national study to assess validity and psychometrics of the short kessler psychological distress scale (k6). *Rehabilitation Counseling Bulletin*, 65(2):140–149.

Watson, K. B., Wiltz, J. L., Nhim, K., Kaufmann, R. B., Thomas, C. W., and Greenlund, K. J. (2025). Trends in multiple chronic conditions among us adults, by life stage, behavioral risk factor surveillance system, 2013-2023. *Preventing Chronic Disease*, 22:E15.

Yu, Z. (2024). Why are older men working more? the role of social security. *Journal of Public Economics*, 231:105071.

Appendix

A Sample selection

Table 7 presents descriptive statistics for the final sample, and Table 8 summarizes how the sample size changes at each stage of selection. I use biennial PSID waves from 2001-2003 and 2007-2021, which together provide an initial sample of 324,106 individual-wave observations. Restricting the analysis to household heads reduces the sample by 233,774 observations. To ensure representativeness, I exclude the Survey of Economic Opportunity (SEO) subsample, which oversamples low-income and minority households, yielding 62,239 observations.

Because the K-6 mental health scale is administered only to household respondents, I further limit the sample to household heads who are also respondents. This condition is met in most cases, but results in the loss of 23,136 observations. To align the sample with the model, I retain only households whose head is between ages 25 and 90. Finally, I drop all observations with non-positive longitudinal weights or missing information on either mental or physical health.

Table 7: Descriptive Statistics, PSID sample.

	# of obs	# of individuals	Average	Std. Dev	Minimum	P25	Median	P75	Maximum
Household income	31,203	6,379	37,418	57209.33	0.00	1,358.79	24,854.39	51,933.42	116,052.64
Labor income	31,203	6,379	29,400	51354.73	0.00	0.00	18,535.62	40,590.95	93,440.01
Wealth	31,203	6,379	290,555	1.20e+06	-17,728.09	4,557.53	56,580.33	237,641.59	1.16e+06
$\mathbb{1}\{m_t = \text{Bad}\}$	31,203	6,379	0.26	0.44	0.00	0.00	0.00	1.00	1.00
K-6 Score	31,203	6,379	3.27	3.95	0.00	0.00	2.00	5.00	12.00
$\mathbb{1}\{p_t = \text{Bad}\}$	31,203	6,379	0.17	0.38	0.00	0.00	0.00	0.00	1.00
General health score	31,203	6,379	2.54	1.04	1.00	2.00	2.00	3.00	4.00
Age	31,203	6,379	51.48	16.76	27.00	37.00	51.00	64.00	81.00
$sex = \text{Women}$	31,203	6,379	0.35	0.48	0.00	0.00	0.00	1.00	1.00
$race = \text{Black}$	28,712	5,806	0.10	0.30	0.00	0.00	0.00	0.00	1.00
Migrant dummy	31,203	6,379	0.07	0.26	0.00	0.00	0.00	0.00	1.00
College dummy	31,203	6,379	0.41	0.49	0.00	0.00	0.00	1.00	1.00

Household income, labor income, and wealth are equivalized and deflated using 2010 U.S. Census CPI

Table 8: Sample selection, PSID waves. Number of observations

Sample selection	Selected out	Selected in
Selected PSID waves		324,106
Heads only	233,774	90,332
Drop SEO sample	28,093	62,239
Respondents only	23,136	39,103
Age between 25 and 90	3,100	36,003
Missing key variables	4,800	31,203

B Retirement Income Calculation

Post-retirement income in the model approximates the structure of the U.S. Social Security system. Individuals retire exogenously at age 65, after which income becomes a deterministic function of pre-retirement earnings:

$$y_{\text{ret}} = f(y_{63}, \overline{y}_{63})$$

where y_{63} is the last realization of labor earnings before retirement and $\overline{y}_{63}$ denotes average earnings in the economy. The mapping follows a replacement-rates schedule consistent with De Nardi et al. (2019).

The replacement schedule is defined as follows: 90 percent of the portion of earnings below 18 percent of average earnings; 32 percent of the portion between 18 and 110 percent of average earnings; and 15 percent of the portion above that threshold. This rule can be summarized as:

$$y_{\text{ret}} = \begin{cases} 0.9y & \text{if } y < 0.18\overline{y}_{63} \\ 0.9 \cdot 0.18\overline{y}_{63} + 0.32 \cdot (y - 0.18\overline{y}_{63}) & \text{if } 0.18\overline{y}_{63} \leqslant y \leqslant 1.1\overline{y}_{63} \\ 0.9 \cdot 0.18\overline{y}_{63} + 0.32 \cdot (1.1\overline{y}_{63} - 0.18\overline{y}_{63}) + 0.15 \cdot (y - 1.1\overline{y}_{63}) & \text{if } y \geqslant 1.1\overline{y}_{63} \end{cases}$$

C First-Step Estimates

C.1 Earning Process

Let $\tilde{y}$ denote "detrended" log earnings, defined as the logarithm of earnings net of the deterministic components associated with age, mental health, and physical health. The identification of the earnings process relies on the following conditions:

$$var(\tilde{y}_{25+s}) = \sigma_\kappa^2 + \sigma_\varepsilon^2 \sum_{t=0}^{s/2} \rho^{2t} + \sigma_\epsilon^2 \tag{18}$$

$$cov(\tilde{y}_{25+m}, \tilde{y}_{25+n}) = \sigma_\kappa^2 + \sigma_\varepsilon^2 \sum_{j=0}^{n/2} \rho^{\frac{(m-n)}{2}+2j} \qquad \text{for } m > n \tag{19}$$

Here, m and n index biennial ages, since PSID waves are collected every two years.

Because the PSID waves used span 2001-2021, the maximum distance between two earnings observations for the same individual is 20 years. Accordingly, the parameters of the AR(1) process are estimated using covariances observed within this horizon. Table 9 reports the estimated variance components and persistence parameter of the AR(1) specification.

Table 9: Estimated Parameters AR(1) process

Parameter	Description	Value
ρ	Persistence of the AR(1) process	0.9504122
σ_κ^2	Variance of individual fixed effect	0.2782947
σ_ε^2	Variance of persistent shock	0.0171284
σ_ϵ^2	Variance of transitory shock	0.3581276

D Estimated health transition probabilities

In this section, I present the estimated probabilities of transitioning to the two remaining health states not discussed in the body of the paper. These transitions capture cases where individuals experience a deterioration in one health dimension while remaining healthy—or improving—in the other.

The top panel of Figure 11 shows the probabilities of transitioning to $h_{t+1} = \{m_{t+1} = \text{Bad}, p_{t+1} = \text{Good}\}$. Across all initial states, this probability declines sharply with age, reflecting that the risk of mental distress diminishes over time, while the risk of physical illness rises. Differences between medium- and high-productivity individuals are minimal, but the low-productivity type stands out: they (i) are more likely to move into this state when initially healthy, (ii) less likely to do so when already physically ill, and (iii) more likely to remain in this state once there.

The bottom panel reports the probability of transitioning to $h_{t+1} = \{m_{t+1} = \text{Good}, p_{t+1} = \text{Bad}\}$. In contrast, these probabilities increase steadily with age, capturing the gradual deterioration of physical health over the life cycle. The upward-sloping profile indicates that physical health shocks become more common at older ages. While differences across productivity groups are negligible when individuals start in poor health in both dimensions, a clear ordering emerges otherwise: the low-productivity individuals face the highest transition probabilities, followed by the medium group, and high-productivity individuals the lowest.

Figure 11: Predicted health transition probabilities by initial health state

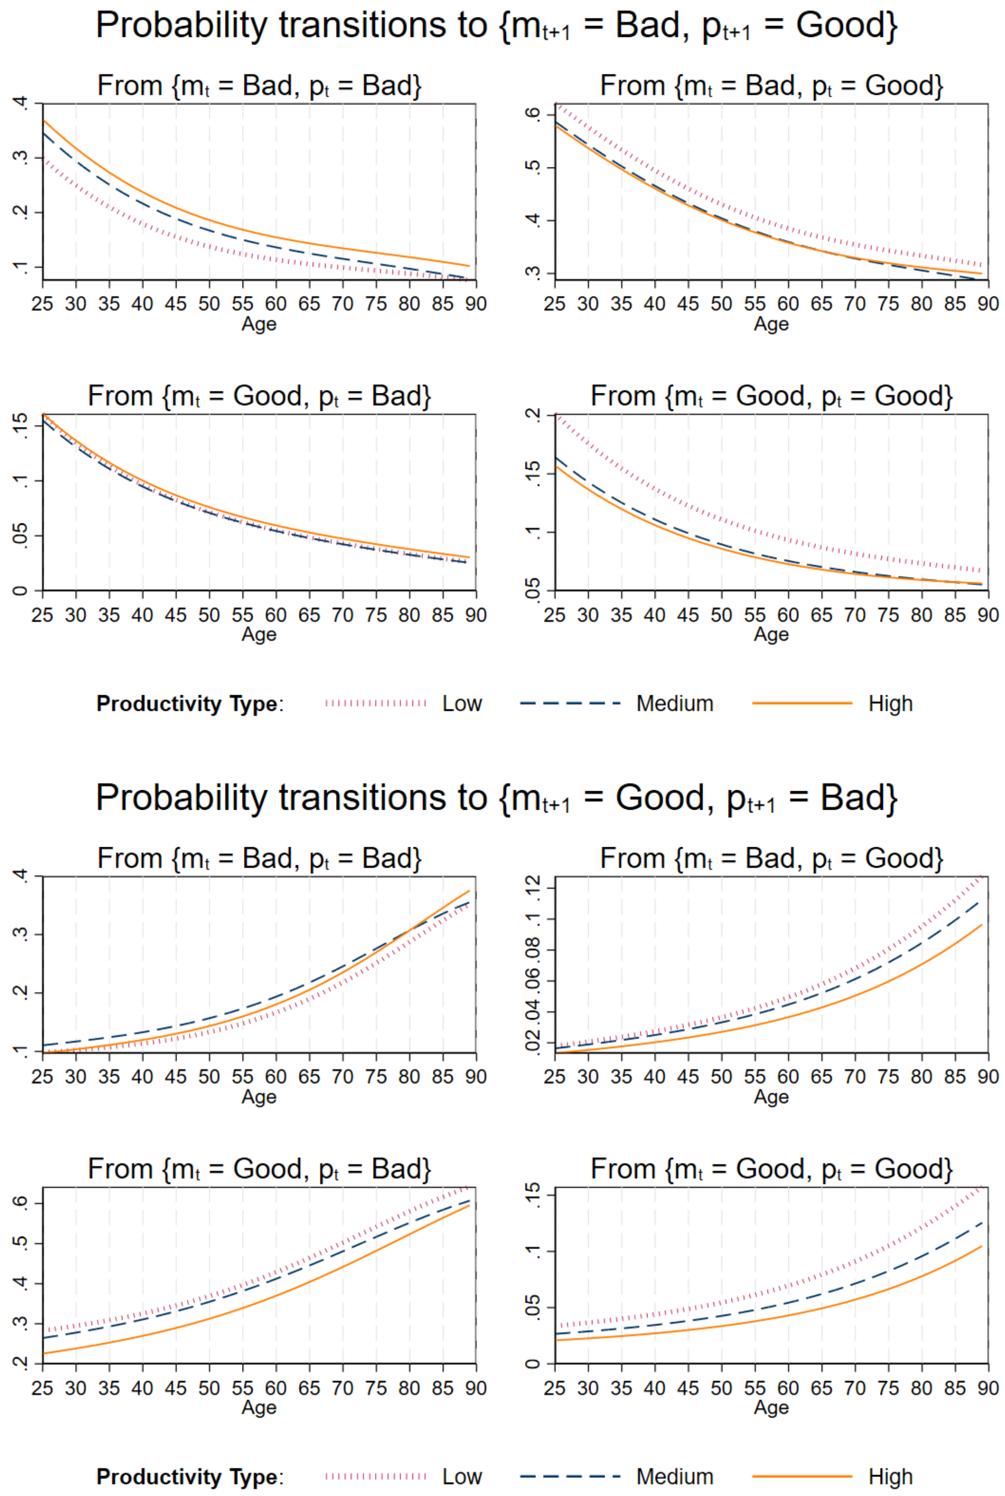

E Model fit for Aggregate Targets

This section evaluates the model's ability to replicate the life-cycle wealth accumulation patterns observed in the PSID. I first assess the aggregate fit, pooling all individuals, and then examine the model's performance when disaggregating by physical health status (good vs. bad). This decomposition enables a direct comparison with existing studies, which typically treat health as a one-dimensional measure based solely on physical condition.

E.1 Aggregate Wealth Moments

Figure 12 compares mean wealth by age group between the data and the simulated economy. The model reproduces the main features of the observed life cycle, capturing both the rapid accumulation of assets before retirement and the moderate decumulation that follows. In particular, it closely matches the sharp rise in wealth leading up to age 65 and the subsequent plateau in asset holdings during retirement.

Figure 12: Empirical vs Model targets: Aggregated Mean Wealth

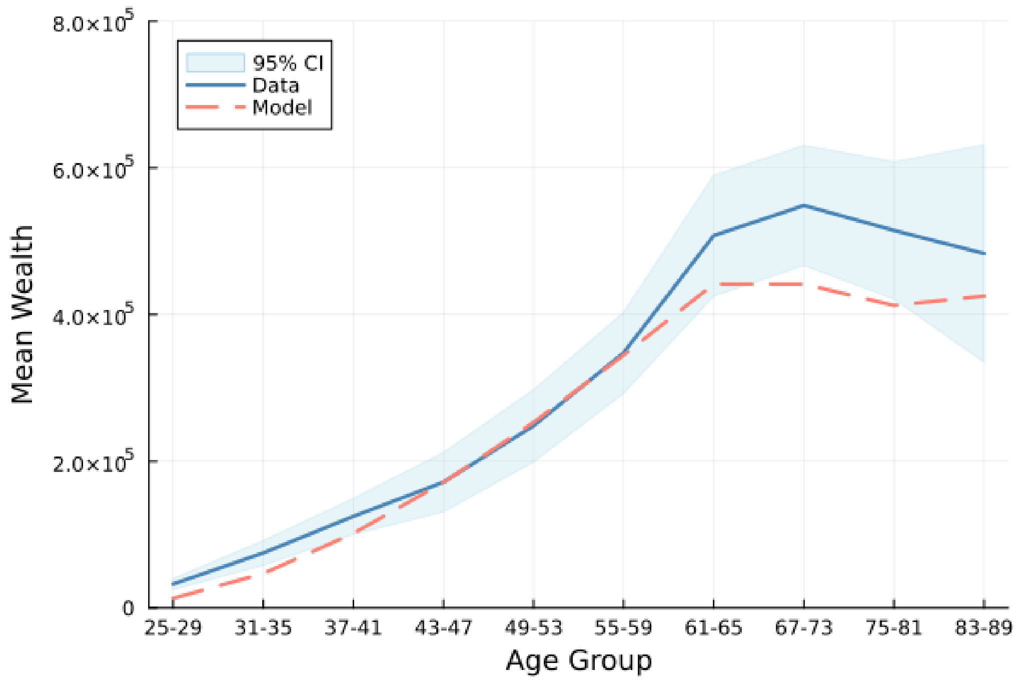

The model slightly underpredicts asset accumulation in the late 60s and early 70s, but the confidence intervals around the empirical moments suggest that these differences are statistically minor. Overall, the aggregate fit suggests that the model accurately captures both the timing and magnitude of wealth accumulation over the life cycle.

E.2 Model Fit by Physical Health

Because most prior studies equate "health" with physical health, I next evaluate the model's fit when grouping individuals solely by physical health status. This aggregation ensures comparability with the existing literature: within each group (physically healthy or unhealthy), individuals may occupy either mental health state.

For individuals in good physical health, the model tracks the data closely across the life cycle. It replicates the steep rise in wealth during the working years and, while slightly underestimating wealth after retirement, accurately matches the level of assets held late in life—consistent with observed bequest behavior. In contrast, for individuals in bad physical health, the model tends to overpredict wealth holdings, particularly from midlife through mid-70s. This overprediction points to potential avenues for refinement in the model.

Figure 13: Empirical vs Model targets: Mean Wealth by Physical Health

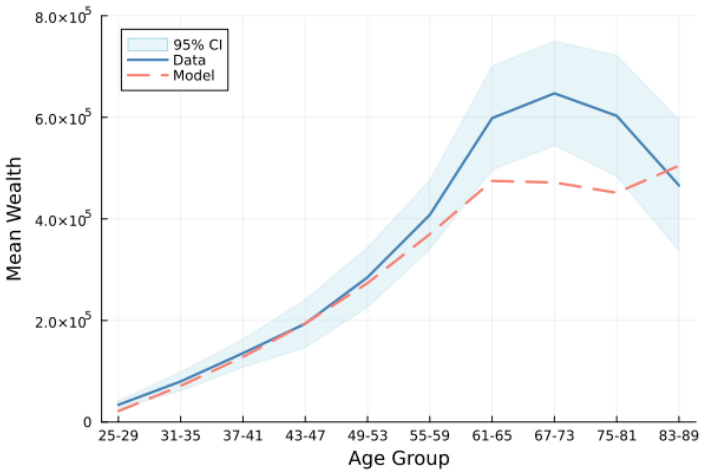

(a) $p_t =$ Good

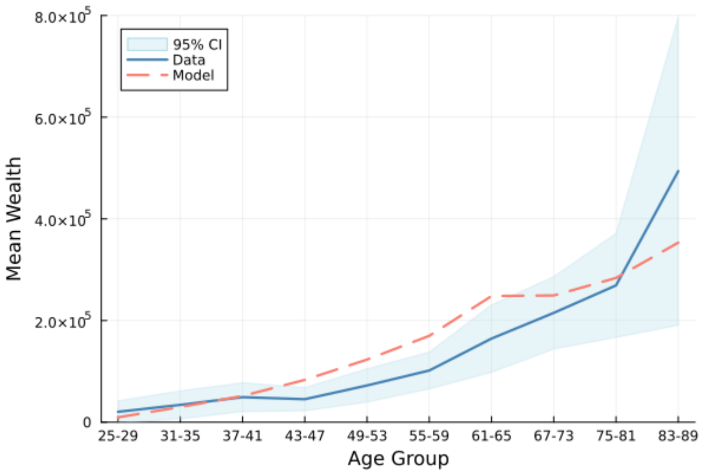

(b) $p_t =$ Bad

Taken together, these results show that the model delivers a strong aggregate fit to wealth dynamics while capturing the heterogeneity between physically healthy and unhealthy individuals. The mild overprediction of late-life wealth among those in poor physical health points to potential refinement in the modeling of health shocks, medical expenses, or mortality risk heterogeneity. Overall, the model's ability to match both the level and trajectory of empirical wealth profiles provides compelling evidence of its internal consistency and supports its use for counterfactual and policy analysis.

F Worsening initial mental health distribution: Results by health state

This section examines how consumption and wealth evolve relative to the baseline when the initial health distribution is worsened, as described in Section 7.

Consistent with the aggregate analysis, worsening initial mental health has limited effects on both consumption and wealth across all health states. Figure 14 shows that the counterfactual produces only modest deviations from the baseline—slightly higher consumption in early adulthood and somewhat lower levels in midlife—yet overall differences remain small.

Similarly, Figure 15 reveals that individuals accumulate marginally less wealth at younger ages under the counterfactual, but these differences vanish as the cohort ages, and the trajectories converge with the baseline. These patterns are consistent across all health states, with no notable divergence between groups.

Overall, the results show that worsening initial mental health primarily shifts the timing of consumption and savings rather than their long-run levels. The effects are modest and uniform across health states, reinforcing that initial conditions influence the intertemporal allocation of resources without substantially altering lifetime economic outcomes.

Figure 14: Counterfactual Analysis: Average Consumption by health state

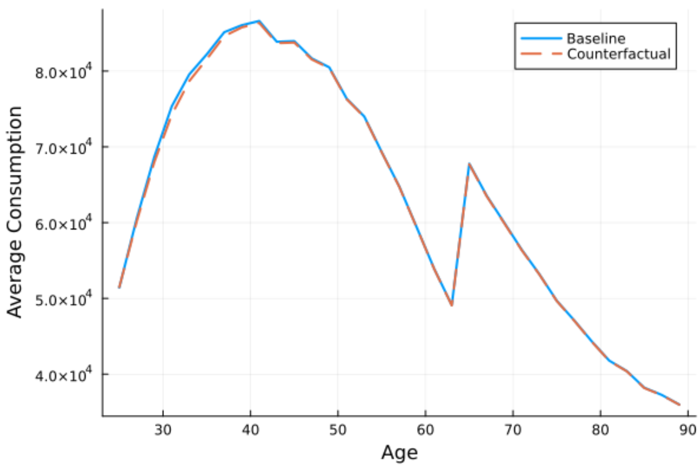

(a) m_t = Good, p_t = Good

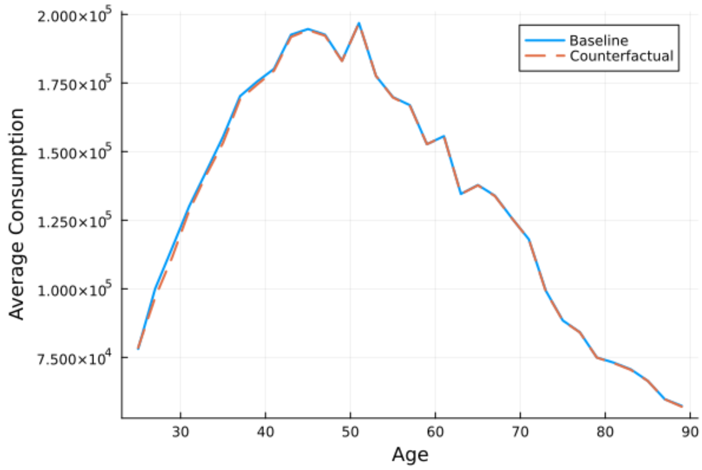

(b) m_t = Bad, p_t = Good

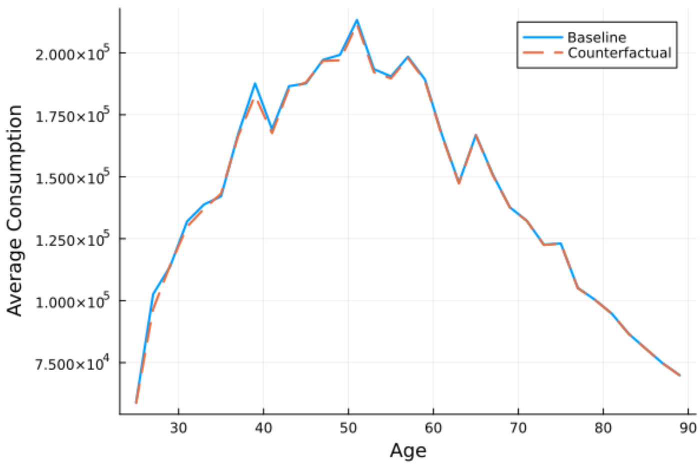

(c) m_t = Good, p_t = Bad

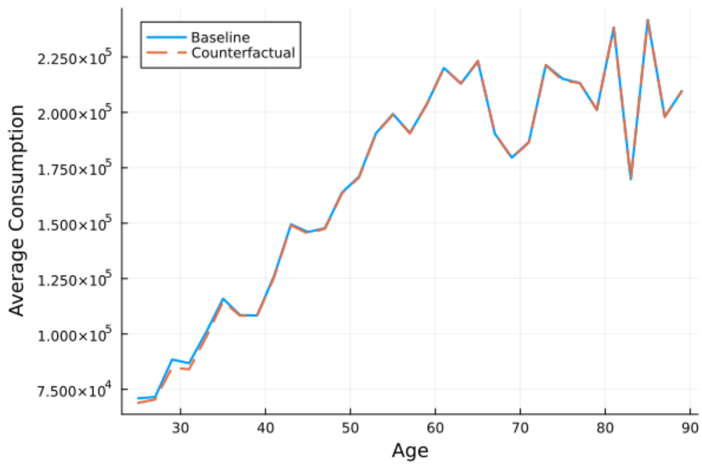

(d) m_t = Bad, p_t = Bad

Figure 15: Counterfactual Analysis: Average Wealth by health state

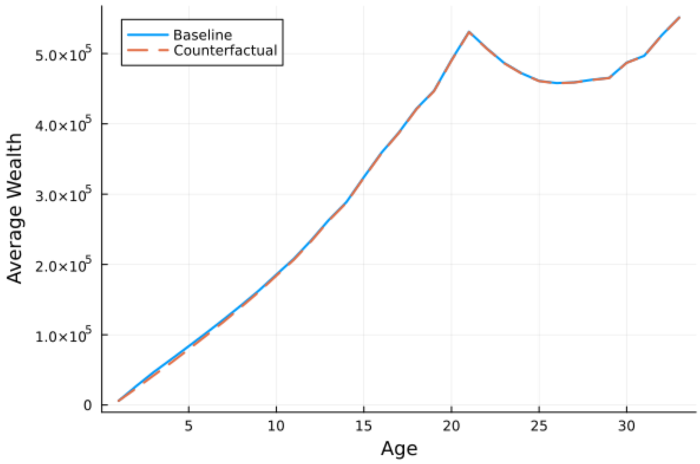

(a) m_t = Good, p_t = Good

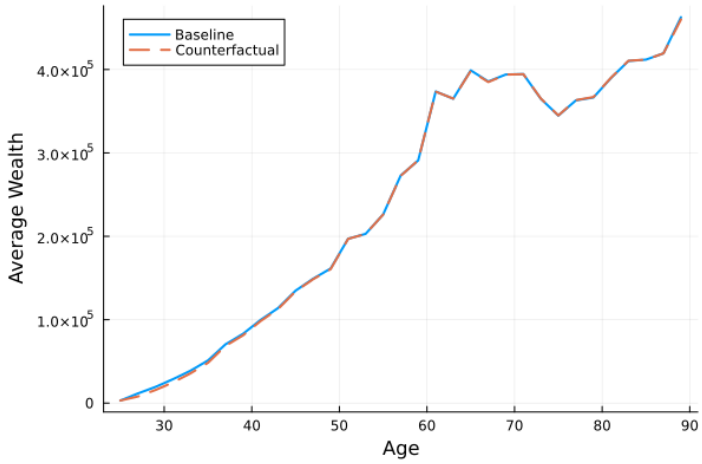

(b) m_t = Bad, p_t = Good

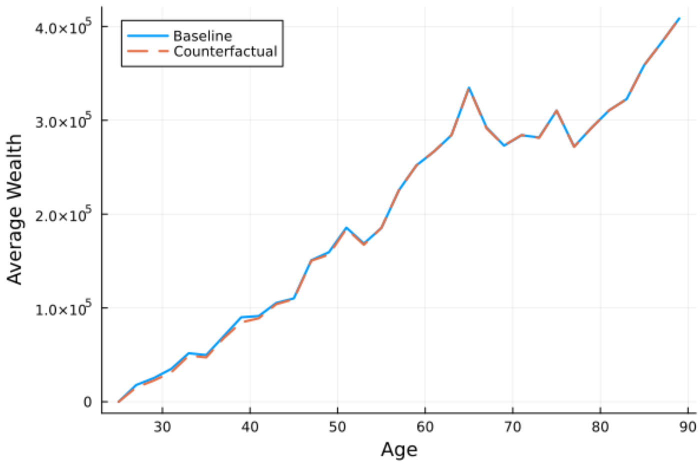

(c) m_t = Good, p_t = Bad

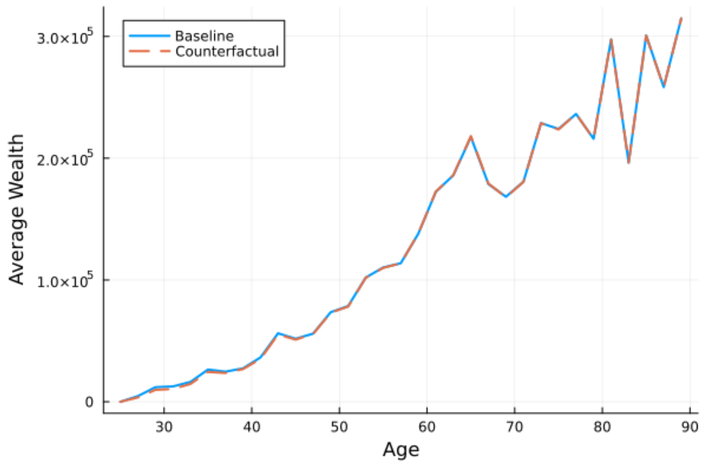

(d) m_t = Bad, p_t = Bad