Optimal Pension Reforms during Demographic Transitions*

Daniel Carroll
Federal Reserve Bank of Cleveland

Sewon Hur
Yonsei University

Selahattin İmrohoroğlu
USC Marshall School of Business

Zhixiu Yu
Louisiana State University

August 5, 2026

[PRELIMINARY AND INCOMPLETE]

Abstract

Demographic transitions threaten pay-as-you-go pension systems. We develop an overlapping generations model with incomplete markets and endogenous retirement to study optimal pension reform during population aging. Using panel data from the Health and Retirement Study, we estimate the elasticity of retirement to financial incentives and find it to be small, limiting the fiscal gains from manipulating retirement incentives. We apply the framework to US Social Security, evaluating three reforms: raising the payroll tax, cutting benefits, and increasing the Normal Retirement Age. We document sharp trade-offs: raising payroll taxes minimizes welfare losses and spreads the burden evenly, while cutting benefits and raising the NRA impose larger costs that fall unevenly across workers and retirees. These findings motivate our Ramsey framework for identifying the optimal, time-varying reform path.

KEYWORDS: retirement, inequality, social security, welfare, aging, reforms.
JEL CLASSIFICATION CODES: E21, H55, J18.

*We thank Minsu Chang, Munseob Lee, and participants at the 2025 SAET conference and the 2026 SED for comments and insightful discussions. The views expressed are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Cleveland or the Federal Reserve System.

1 Introduction

Across the OECD, rising life expectancy and declining fertility are increasing old-age dependency ratios and threatening the solvency of pay-as-you-go pension systems. In the United States, the Social Security system faces long-run fiscal imbalances as the ratio of retirees to workers rises. Policymakers must balance fiscal sustainability against the twin goals of minimizing distortions to labor supply and savings and ensuring adequate retirement consumption. How should the burden of adjustment be allocated—through higher taxes on workers, reduced benefits for retirees, or later retirement? And how do these policy choices interact with the decision to retire, which is itself responsive to economic incentives?

We address these questions using a quantitative overlapping generations model with incomplete markets and endogenous retirement. A distinguishing feature of our approach is that the retirement decision is disciplined by micro-data. Using panel data from the Health and Retirement Study (HRS), we estimate the elasticity of retirement with respect to financial incentives and find it to be small: a one-unit increase in log wealth raises the probability of retirement by approximately 0.2 percentage points. This low elasticity has a key policy implication—it limits the fiscal gains from manipulating retirement incentives alone and implies that reforms must rely on a combination of instruments.

We evaluate three reform scenarios for US Social Security: increasing the payroll tax, cutting benefits, and raising the Normal Retirement Age (NRA). Each involves starkly different macroeconomic and distributional consequences. Raising the payroll tax minimizes aggregate welfare losses and spreads the burden evenly across workers and retirees. Cutting benefits hits retirees hardest. Raising the NRA shifts the burden to workers and works primarily as an implicit benefit cut through the PIA adjustment, since the low retirement elasticity means average retirement age barely changes. These sharp trade-offs motivate our development of a Ramsey framework for identifying the optimal, time-varying reform path.

Our paper contributes to a large literature that uses quantitative general equilibrium models with overlapping generations to study public pension systems. Following the seminal work of Auerbach and Kotlikoff (1987), Imrohoroglu, İmrohoroğlu, and Joines (1995) introduced uncertain lifetimes and borrowing constraints to evaluate the welfare effects of unfunded social security. Huang, İmrohoroğlu, and Sargent (1997) and Nishiyama and Smetters (2007) extended this framework to compute equilibrium transition paths, Conesa and Krueger (1999) showed that idiosyncratic uncertainty significantly reduces political support for privatization, and De Nardi, İmrohoroğlu, and Sargent (1999) and Attanasio, Kitao, and Violante (2007) incorporated realistic demographic projections to quantify transitional welfare costs. More recently, Huggett and Parra (2010) analyze optimal piecemeal reforms, Krueger and Kubler (2006) show that Pareto-improving reforms are feasible under incomplete markets, Hosseini and Shourideh (2019) characterize Pareto optimal reforms in an OLG economy, and Makarski, Tyrowicz, and Komada (2024) highlight the importance of fiscal closure rules. McGrattan and Prescott (2017) demonstrate

large welfare gains from transitioning to a saving-for-retirement system when productive capital is properly measured.

Our work is most closely related to studies that integrate endogenous retirement into the evaluation of pension policy. Building on French (2005), İmrohoroğlu and Kitao (2012) demonstrate that endogenous labor supply at both the intensive and extensive margins significantly alters predicted solvency gains. Díaz-Giménez and Díaz-Saavedra (2009) show that raising retirement ages can restore sustainability in Spain, Kitao (2014) compares reform options for the US, and Rogerson and Wallenius (2013) and Yu (2024) highlight how nonconvexities and Social Security rules drive labor supply adjustments at older ages. On the normative side, Shourideh and Troshkin (2017) find that optimal systems should reward later retirement regardless of the efficient retirement age profile, Ndiaye and Yu (2025) show that flexible policies outperform rigid increases in statutory retirement ages, and Jones and Li (2023) and Bairoliya and McKiernan (2021) explore how heterogeneous mortality risks shape welfare-maximizing rules.

Relative to this literature, we make three contributions. First, we estimate the retirement elasticity from micro-data and use these estimates to discipline the model's endogenous retirement margin—a step that is critical for credibly evaluating the fiscal efficacy of retirement-age reforms. Second, we provide a comparative analysis of reform instruments for the US and decompose transition welfare into price and policy channels. Third, we develop a Ramsey framework for solving for the welfare-maximizing time-varying reform path, building on the approach of Dyrda and Pedroni (2023).

The remainder of the paper is structured as follows. Section 2 presents our general equilibrium overlapping generations model with incomplete markets and endogenous retirement. Section 3 calibrates the model to the United States and presents policy experiments under several reform scenarios. Section 5 sets up the Ramsey framework for determining optimal policies. Finally, Section 6 concludes.

2 Model

The economy is populated by overlapping generations of households who live for at most J periods. Age is denoted by $j = 1, \ldots, J$. A newborn enters the economy with assets a, an idiosyncratic labor-productivity state z, permanent education type $s \in \mathcal{S} \equiv \{c, n\}$, health status $h \in \mathcal{H} \equiv \{g, b\}$, and accumulated pension earnings e. The type c denotes college and n denotes non-college. The health state g denotes good health and b denotes bad health. Households face incomplete markets, stochastic productivity, stochastic health, mortality risk, and medical-expenditure risk. They choose consumption, saving, labor supply, and whether to claim pension benefits when eligible.

2.1 Demographics, Types, and Health

A measure one of newborns enters in each period, and cohort size grows at the exogenous rate n_t. At birth, households draw a permanent education type $s \in \{c, n\}$ with probabilities π_s, where $\pi_c + \pi_n = 1$. This type is fixed over the life cycle and affects earnings profiles, health transitions, medical expenditures, survival probabilities, and initial distributions.

Health evolves in every period. Conditional on age j, education s, and current health h, next period's health is drawn from

$$\Pi_j^h(h' \mid h, s) \equiv \Pr(h_{j+1} = h' \mid h_j = h, s), \qquad h, h' \in \{g, b\}. \tag{1}$$

Survival from age j to age $j+1$ is denoted by $\psi_j(s, h)$, with $\psi_J(s, h) = 0$. Health therefore matters for expected longevity and, through that channel, for saving, claiming, and retirement decisions.

Each period, after the health state is realized, the household draws an out-of-pocket medical expenditure shock

$$m \sim \Pi_j^m(\cdot \mid s, h), \tag{2}$$

where m is net of any public or private insurance payments. The distribution is age-, education-, and health-specific. Bad health is allowed to increase both the mean and the dispersion of m. The timing follows Imrohoroglu and Kitao (2012): health and the medical expenditure shock are realized before the household chooses consumption, labor supply, and savings. The medical expenditure must be paid within the period and directly reduces disposable resources.

2.2 Endowments and Preferences

A working-age household is endowed with one unit of time. If it works $\ell \in [0, 1]$ hours, its effective labor supply is

$$n_j(s, h, z, \ell) = \eta_j^s \, \omega_h \, z \, \ell, \tag{3}$$

where η_j^s is the education-specific deterministic age-efficiency profile, ω_h is the health-specific productivity factor with $\omega_g \geq \omega_b > 0$, and z is the idiosyncratic labor-productivity shock. Productivity follows a finite-state Markov chain

$$\Pi_j^z(z' \mid z, s) \equiv \Pr(z_{j+1} = z' \mid z_j = z, s). \tag{4}$$

Flow utility is

$$u(c, \ell) = \frac{c^{1-\sigma} - 1}{1 - \sigma} - \theta_\ell \frac{\ell^{1+\chi}}{1 + \chi}, \tag{5}$$

with log when $\sigma = 1$. Retired households set $\ell = 0$. Households discount the future at rate β. If a household dies with end-of-period assets a', it receives warm-glow utility

$$\Omega(a') = \theta_b \frac{(a' + \kappa_b)^{1-\gamma_b} - 1}{1 - \gamma_b}, \tag{6}$$

where θ_b governs the strength of the bequest motive and κ_b is a shifter.

2.3 Technology and Factor Prices

A representative firm operates a constant-returns technology

$$Y_t = F(K_t, N_t) = Z_t K_t^{\alpha} N_t^{1-\alpha}, \tag{7}$$

where K_t is aggregate capital and N_t is aggregate effective labor. Capital depreciates at rate δ. Competitive factor prices satisfy

$$r_t^K = \alpha Z_t \left(\frac{N_t}{K_t}\right)^{1-\alpha} - \delta, \tag{8}$$

$$w_t = (1-\alpha) Z_t \left(\frac{K_t}{N_t}\right)^{\alpha}. \tag{9}$$

2.4 Assets, Government Debt, and Returns

Households save in a one-period riskless asset $a' \geq \underline{a}$. Aggregate household assets are claims on physical capital and government debt. Let D_t denote government debt and K_t physical capital. Aggregate assets satisfy

$$A_t = K_t + D_t. \tag{10}$$

For notational simplicity, household budget constraints use the common net return r_t. In a stationary equilibrium with no arbitrage between government debt and capital claims, $r_t = r_t^K = r_t^D$. If the model is extended to allow wedges between the two assets, r_t should be interpreted as the after-tax portfolio return relevant for household saving.

2.5 Government, Taxes, and the Pension System

The government purchases G_t, issues debt D_t, levies taxes, provides a consumption floor, and operates a pay-as-you-go pension system. Labor income is

$$y_t = w_t \eta_j^s \omega_h z \ell. \tag{11}$$

Workers pay a payroll tax τ_t^{ss} up to a taxable maximum $\bar{y}_t^{ss}$ and pay income taxes according to $T_t(y_t)$. Pension benefits are a function of accumulated pension earnings e and claiming age. Let $b_t(e, j^R)$ denote the annual benefit for a household that claims at age j^R, incorporating the benefit formula, early-claiming reductions, and delayed-retirement credits. The state variable e evolves with covered earnings while the household works,

$$e' = \mathcal{E}_j(e, y_t), \tag{12}$$

and is fixed after claiming.

The model includes a minimum consumption level, denoted by $\underline{c} > 0$. The floor is implemented after the health and medical shocks are realized and before choices are made.

Let x denote cash on hand net of taxes and medical expenses but before government transfers. The transfer $tr_t(x)$ is the smallest nonnegative amount that permits the household to choose at least $\underline{c}$ while satisfying the borrowing constraint:

$$tr_t(x) = \max\{0, \underline{c} + \underline{a} - x\}. \tag{13}$$

Equivalently, the household's feasible set satisfies $c \geq \underline{c}$ and $a' \geq \underline{a}$. This timing captures the role of Medicaid and other safety-net programs that absorb medical-expense shocks that would otherwise push consumption below the floor.

2.6 Timing

Within each period, events occur in the following order:

1. The household enters the period with state (j, a, z, e, s, h, q), where $q \in \{0, 1\}$ indicates whether pension benefits have already been claimed.
2. The current health state h is observed and the out-of-pocket medical expenditure m is drawn from $\Pi_j^m(\cdot \mid s, h)$.
3. The household chooses consumption c, labor supply ℓ, next-period assets a', and, if eligible and not yet claimed, whether to claim pension benefits.
4. Production takes place, taxes are paid, pension benefits and transfers are received, and medical expenses are paid.
5. Mortality is realized. Survivors draw next period's health and productivity states. Deceased households receive warm-glow utility from bequests, and their remaining assets are distributed according to the estate and transfer rules.

2.7 Household Problem

Let $q \in \{0, 1\}$ denote claiming status, where $q = 0$ means not yet claimed and $q = 1$ means claimed. Let $V_{j,t}(a, z, e, s, h, q)$ be the value function at age j in period t before the medical shock is drawn. Conditional on m, the household solves

$$\begin{aligned} V_{j,t}(a, z, e, s, h, q) = \mathbf{E}_{m|j,s,h} \Big\{ & \max_{c,\ell,a',q'} \; u(c, \ell) \\ & + \beta\psi_j(s, h)\mathbf{E}_{z',h'} V_{j+1,t+1}(a', z', e', s, h', q') + \beta[1 - \psi_j(s, h)]\Omega(a')\Big\} \end{aligned} \tag{14}$$

$$\begin{aligned} \text{s.t.} \quad & c + a' + m_j(h) \leq (1 + r_t)a + y_t - T_t(y_t) - \tau_t^{ss} \min\{y_t, \bar{y}_t^{ss}\} + q' b_t(e, j^R) + tr_t(x), \\ & y_t = w_t \eta_j^s \omega_h z \ell, \\ & e' = \begin{cases} \mathcal{E}_j(e, y_t), & q' = 0, \\ e, & q' = 1, \end{cases} \\ & c \geq \underline{c}, \qquad a' \geq \underline{a}, \qquad \ell \in [0, 1]. \end{aligned}$$

The claiming choice is restricted by eligibility. If $j < j_E$, then $q' = 0$. If $q = 1$, then $q' = 1$. If $j \geq j_E$ and $q = 0$, the household may set $q' = 1$ and starts receiving the permanently adjusted benefit $b_t(e, j)$ in the current period. Retired households are the special case with $q = 1$; workers are households with $q = 0$.

2.8 Recursive Equilibrium

A recursive competitive equilibrium is a sequence of household value functions and policy functions, distributions of households μ_t, aggregate quantities $\{K_t, N_t, A_t, D_t, Y_t\}_{t\geq 0}$, prices $\{r_t, w_t\}_{t\geq 0}$, government policies $\{T_t, \tau_t^{ss}, b_t, tr_t, G_t, D_{t+1}\}_{t\geq 0}$, and laws of motion for the distribution such that:

1. Given prices and government policies, household policy functions solve the dynamic program in (14) for every state (j, a, z, e, s, h, q).

2. Firms maximize profits, and factor prices equal marginal products.

3. Markets clear:

$$K_t + D_t = \int a \, d\mu_t, \tag{15}$$

$$N_t = \int \eta_j^s \omega_h z \ell_{j,t}(a, z, e, s, h, q, m) \, d\mu_t, \tag{16}$$

$$Y_t = C_t + K_{t+1} - (1 - \delta)K_t + G_t + M_t, \tag{17}$$

   where C_t is aggregate consumption and M_t is aggregate out-of-pocket medical expenditure.

4. The government budget constraint holds each period:

$$G_t + B_t + TR_t + (1 + r_t^D)D_t = T_t^{inc} + T_t^{ss} + D_{t+1}, \tag{18}$$

   where B_t denotes aggregate pension benefits, TR_t aggregate consumption-floor transfers, T_t^{inc} income-tax revenue, and T_t^{ss} payroll-tax revenue.

5. The pension system budget constraint is satisfied under the chosen closure rule. For example, under period-by-period balance,

$$B_t = T_t^{ss}. \tag{19}$$

6. The distribution μ_{t+1} is generated from μ_t, household policy functions, survival probabilities, health transitions, productivity transitions, education-type shares, and the rules governing bequests and newborn initial conditions.

A stationary equilibrium is a recursive competitive equilibrium in which prices, policies, age-normalized distributions, and aggregate quantities are constant over time.

3 Calibration

The calibration follows a two-step structure. First, parameters that can be identified outside the model are estimated directly from micro data or set using institutional rules. Second, the remaining preference and fiscal parameters are calibrated jointly so that the stationary equilibrium matches key life-cycle and aggregate moments. Table 4 summarizes the full strategy.

3.1 Health, Mortality, and Medical Spending

In each time period, individuals face uncertainty in health, $h_t \in \{0, 1\}$, with 0 being in good health, and 1 being in bad health. Health in the next period, $h_{t+1} \in \{0, 1\}$, depends on the individual's education, current health, and age, and evolves according to the Markov chain between two states, with an age-dependent Markov transition matrix. A typical element of the health transition matrix at age t with education level c is given by:

$$\pi_{j,i,t+1,c} = Pr(h_{t+1} = j | h_t = i, t + 1, c), \qquad i, j \in \{0, 1\} \tag{20}$$

The lifespan is uncertain. The parameter s_{t+1} denotes the probability that an individual is alive at age $t + 1$ conditional on being alive at age t. The survival probability depends on age and previous health status, as:

$$s_{t+1} = S(h_t, t + 1, c), \qquad h_t \in \{0, 1\} \tag{21}$$

Because individuals live up to a maximum age J, $s_{J+1} = 0$ for any h and c.

Let m_t denote the out-of-pocket medical expenditure at age t, which depends on age and health status,

$$m_t = M(h_t, t, c). \tag{22}$$

The m_t is defined as the individual's total medical expenditure net of medical coverage provided by insurance, such as Medicaid, Medicare, etc. Age- and health-dependent medical expenditure are intended to capture the fact that people spend more on out-of-pocket medical services as they age or health status gets worse.

3.2 Wages

The logarithm of hourly wages at age t, $\ln w_t$, is a function of health status and age, $W(h_t, t)$, and an autoregressive component of wages, ω_t, as follows:

$$\ln w_t = W(h_t, t, c) + \omega_t \tag{23}$$

$$\omega_t = \rho\omega_{t-1} + \eta_t, \quad \eta_t \sim N(0, \sigma_\rho^2) \tag{24}$$

The autoregressive component ω_t has the correlation coefficient, ρ, and a normally distributed innovation, η.

3.3 First-Step Estimation

We primarily use two data sets for estimating the parameters in the first step: the Panel Study of Income Dynamics (PSID) and the Medical Expenditure Panel Survey (MEPS). Table 1 summarizes the first-step parameters, their values, and their sources.

Table 1: First-Step Parameters Summary

Parameter	Description	Value	Source
Wage-Related Parameters			
$W(\cdot)$	Deterministic wages	text	PSID
ρ	Autoregressive coefficient	0.984	PSID
σ_ρ^2	Variance of innovation	0.012	PSID
Mortality and Health Transitions			
s_{t+1}	Survival probabilities	text	PSID
$\pi_{h_t,h_{t-1},t}$	Health transitions	text	PSID
$m(\cdot)$	Out-of-pocket medical expenses	text	MEPS

3.3.1 Data

Our principal data source is the 1968-2015 PSID. Additionally, we use data from MEPS to estimate the profiles of medical expenditures. we collect the 1999–2012 waves and drop the observations with missing values of relevant variables, e.g., age, medical spending, or health insurance.

3.3.2 Health Transitions

Health status is measured based on the self-reported health questions. Respondents rank their health as *(1) excellent, (2) very good, (3) good, (4) fair, or (5) poor*. Individuals whose health rank falls in the first three categories in all three waves are referred to people in good health ($h_t = 0$) and the rest are referred to as people in bad health ($h_t = 1$).

Since there are two health states, $h_t = \{0, 1\}$, we have four transition patterns for each education type c. We estimate the evolution of health by running a probit regression for an indicator $(h_t = j|h_{t-1} = i)$ on age dummies using the sample with $(h_{t-1} = i)$. The predicted values of these probit regressions are the estimates of $Pr(h_t = j|h_{t-1} = i)$, $i, j \in \{0, 1\}$. We linearly interpolate the probabilities of transitioning in three states across ages so that the transition matrices changes smoothly over the life cycle. Figure 1 displays the age-specific health transition out of good health status $Pr(h_t|h_{t-1} = 0)$ and bad health status $Pr(h_t|h_{t-1} = 1)$.

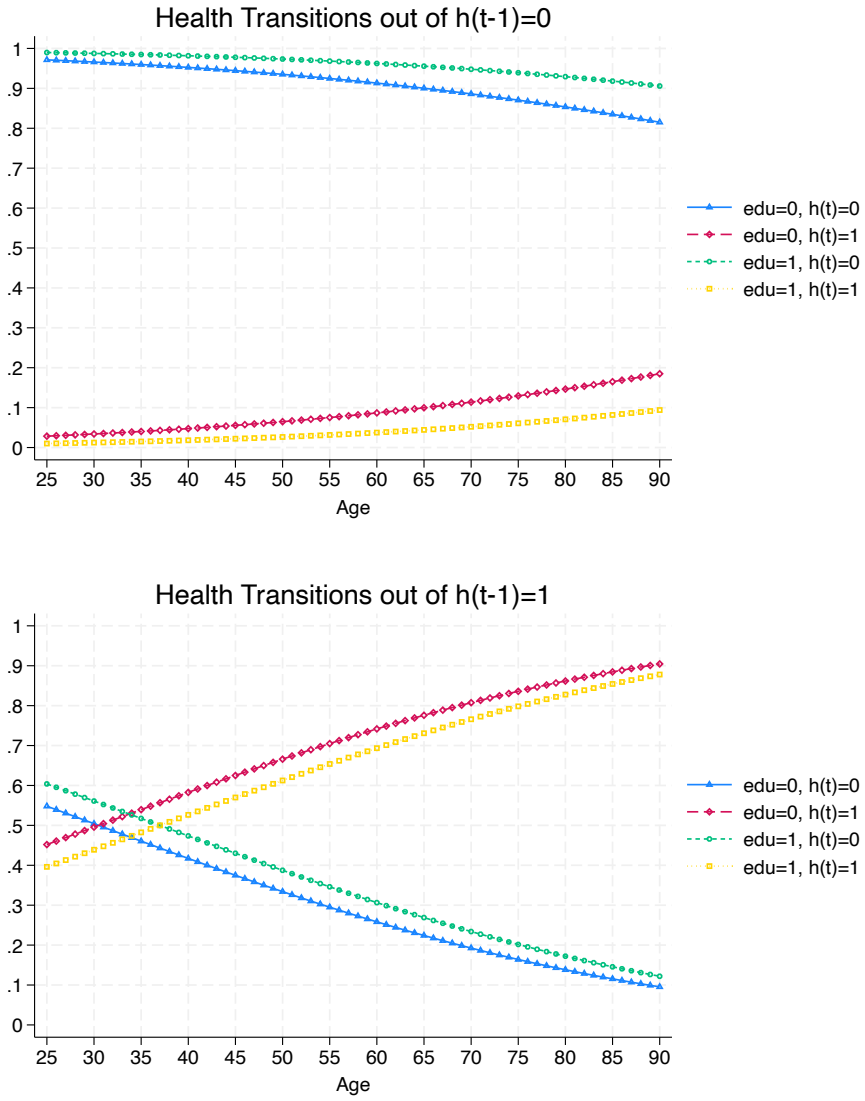


Figure 1: Health transitions and aging

3.3.3 Survival Probabilities

For each education type, age- and health-dependent survival probabilities are estimated by running a logistic regression for the indicator of survival on age polynomial, previous health status, education level, and interaction between age and health, using data from the PSID. The estimated coefficients are reported in Table 2. Figure 2 displays the estimated mortality rates by age and health status. As French(2005) discussed, the PSID underestimates mortality rates by 25%. We adjusted my estimated mortality rates by multiplying 1.25 for the model estimation.

Table 2: Logit Estimates of Survival

	Survival
Age squared	-0.000638^{***}
	(0.0000390)
Bad health $(t-1)$	-2.322^{***}
	(0.455)
College	0.524^{***}
	(0.133)
Bad health $(t-1) \times$ age	0.0134^{**}
	(0.00674)
Constant	7.991^{***}
	(0.186)
Observations	94,126
Log likelihood	$-2{,}879.0$

Notes: Logit estimates; the dependent variable is an indicator for surviving to $t+1$. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

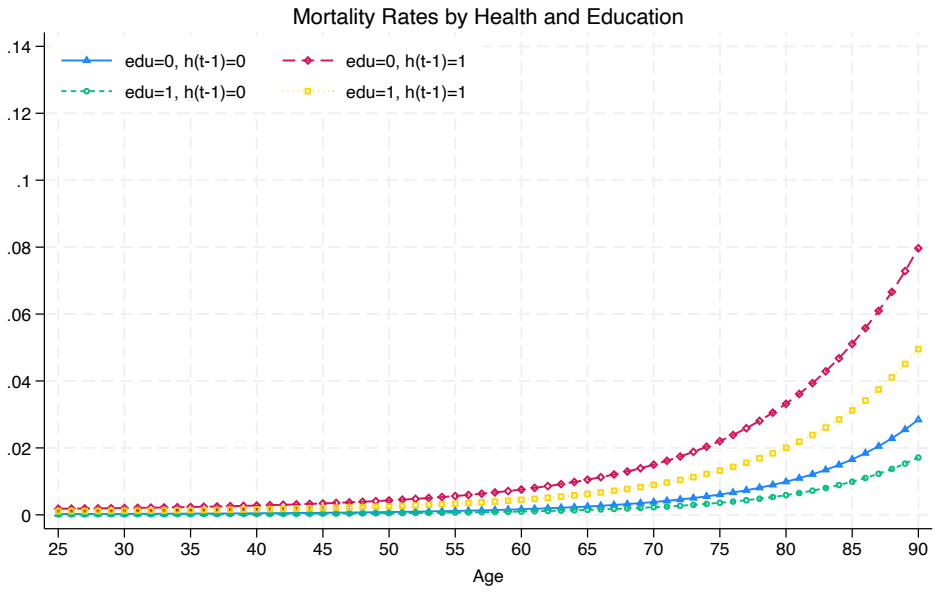


Figure 2: Mortality Rates by Health Status and Education

3.3.4 Out-of-Pocket Medical Expenses

The out-of-pocket medical expenses are computed as total medical expenditure net of the amount covered by health insurance programs, using data from the MEPS. Health status is measured based on self-perceived health rank. We estimate the out-of-pocket expenditure in the following steps.

First, we estimate the profiles of total medical expenses by running a weighted regression on age dummies and year dummies, separately for individuals in each health status. Then we estimate the profiles of insurance coverage by regressing them on age dummies and year dummies, separately by health status. For working-age individuals (age 25-64), the medical coverage is computed as the amount paid by insurance programs. In contrast, for individuals who are 65 or older, the coverage is calculated as the expenses paid by Medicare. Next, we compute the out-of-pocket expenditure as the product of estimated coinsurance rates and total medical expenses at each age and health status. Following De Nardi et al (2018), the estimated medical expenses are multiplied by 1.60 for people younger than 65 years old and by 1.90 for people who are 65 or older to make medical spending consistent with the aggregate medical expenditures from the National Health Expenditure Account (NHEA). Last, we smooth the estimated profiles by regressing them on a quadratic function of age.

Figure 3 displays the smoothed life-cycle profiles of the out-of-pocket medical expenses by health status and education type.

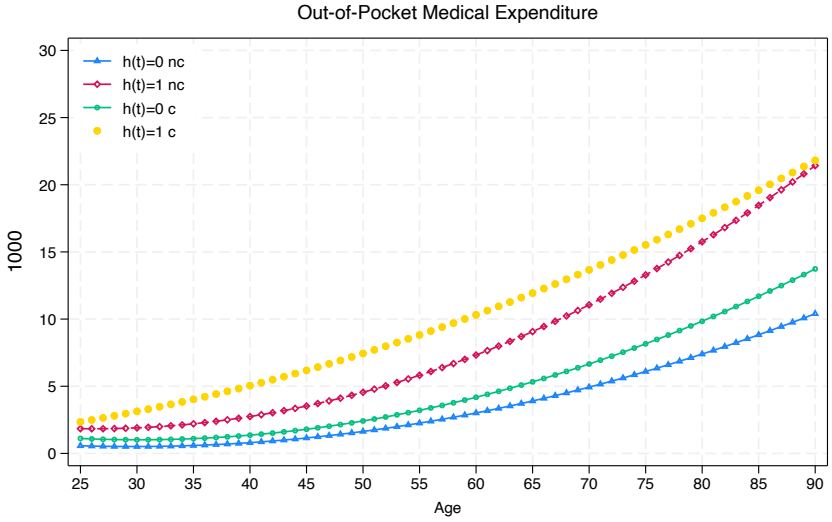


Figure 3: Out-of-Pocket Medical Expenditure by Health Status and Education

3.3.5 Hourly Wages

Hourly wages are computed as annual earnings divided by annual hours worked, using data from the PSID. Respondents in year t report their annual earnings and hours worked in year $t-1$. We drop the observations with hourly wages below \$6.50 or above \$250 (in 2016 dollars) to control the minimum wage and high wage outliers. Since wages are only observed for labor market participants in the data, to adjust selection bias in observed wages, we estimate the age and health-specific hourly wage profiles using the Heckman selection model. Specifically, in the first step, we estimate the selection equation (labor force participation) by running a probit regression using all observations from the PSID, and an inverse Mill's ratio is generated. In the second step, we estimate hourly wage by running a regression on age polynomials, the intersection of health and age polynomials, and the inverse Mill's ratio obtained from the first step. The estimated coefficients of the two-stage procedure are reported in Table 3.

Individuals' work decisions are strongly affected by their life-cycle wage profiles. As described in the model section, the deterministic wage profiles $W(h_t, t, c)$ depend on the individual's age, education, and health.

Figure 4 displays smoothed wage profiles by health status and education. We es-

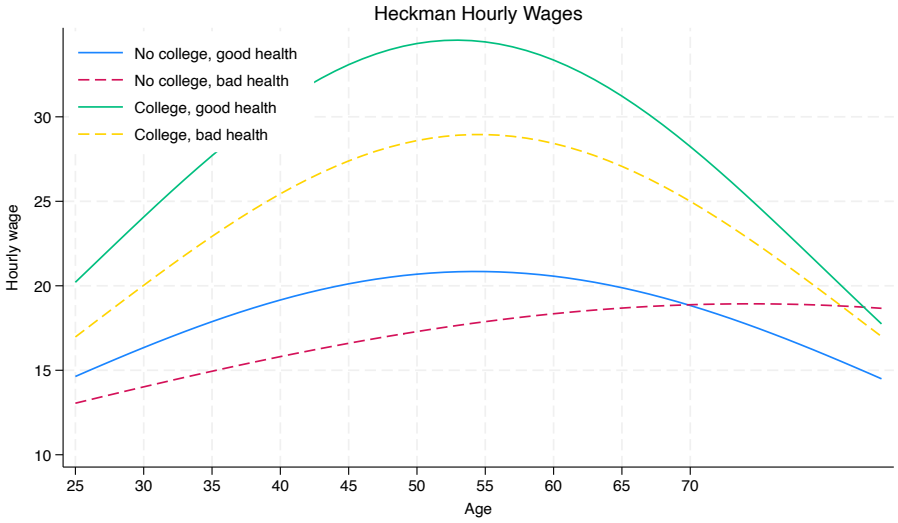


Figure 4: Hourly Wages by Health and Education

timate the stochastic components of hourly wages – the autoregressive coefficient and variance of wage shocks (ρ, σ_ρ^2), using the wage residuals from the above steps. Following the procedure described in Borella et al (2019), we limit the age range between 30 and 75 and drop the highest 0.5% residuals to avoid large outliers that inflate the variance. The estimation process is performed by Maximum Likelihood and standard Kalman Filter recursions. The estimated results for (ρ, σ_ρ^2) in wages are (0.984, 0.012). The estimated stochastic components of wages and deterministic age- and health-specific profiles are

used in Equation (23) to simulate wages and fed into the model estimation.

Table 3: Heckman Selection Estimates of Hourly Wages

Wage equation (log hourly wage)		Selection equation (participation)	
Non-college, bad health	0.275**	Age	−0.00825
	(0.125)		(0.0109)
College, good health	−0.207***	Age^2	0.00139***
	(0.0612)		(0.000240)
College, bad health	−0.278	Age^3	−0.0000213***
	(0.329)		(0.00000170)
Age	0.0448***	Bad health	−0.654***
	(0.00269)		(0.0125)
Non-college, bad health × age	−0.0221***	College	0.205***
	(0.00548)		(0.00971)
College, good health × age	0.0281***	Family size	−0.0199***
	(0.00288)		(0.00176)
College, bad health × age	0.0221	Birth year	0.00412***
	(0.0141)		(0.000614)
Age^2	−0.000412***	Constant	−7.549***
	(0.0000344)		(1.227)
Non-college, bad health × age^2	0.000259***	Inverse Mills ratio (λ)	−0.135***
	(0.0000599)		(0.0396)
College, good health × age^2	−0.000276***		
	(0.0000328)		
College, bad health × age^2	−0.000201		
	(0.000149)		
Constant	1.471***		
	(0.0612)		
Observations	116,311		

Notes: Heckman two-step estimates. The omitted category in the wage equation is non-college workers in good health. The selection equation is a probit for labor force participation. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Externally estimated processes

Education types. The population is split into college and non-college permanent types, $s \in \{c, n\}$. The college share π_c is measured from the relevant cohort in the data. Initial assets, productivity, pension earnings, and health are allowed to differ by education type.

Health and mortality. Health is binary, $h \in \{g, b\}$. Following the empirical strategy in Bairoliya and McKiernan (2026), health status can be constructed from self-reported health by grouping excellent, very good, and good health into g, and fair or poor health into b. Age- and education-specific health transition matrices $\Pi_j^h(h' \mid h, s)$ are estimated

from longitudinal health data. Survival probabilities $\psi_j(s,h)$ are estimated by age, education, and health, and then aligned with life-table mortality by cohort.

Medical expenditures. Out-of-pocket medical expenditures are estimated by age, education, and health status. The model uses the full conditional distribution $\Pi_j^m(m \mid s,h)$ rather than only the mean, since the right tail of medical spending is important for precautionary saving and for the value of the consumption floor. The empirical measure should be net of insurance payments and should include premiums when they are paid out of pocket.

Labor productivity. The deterministic component of earnings is estimated separately by education type, yielding η_j^c and η_j^n. The idiosyncratic component z follows an education-specific Markov process, with persistence and innovation variance estimated from panel earnings data. Health affects contemporaneous productivity through ω_h.

Social Security rules. Payroll taxes, the taxable maximum, the PIA formula, bend points, early-claiming reductions, delayed-retirement credits, and the normal retirement age are set to match U.S. Social Security rules for the cohort under study. The state variable e summarizes average indexed earnings relevant for benefits.

Consumption floor. The minimum consumption level $\underline{c}$ follows the timing in Imrohoroglu and Kitao (2012). The transfer is computed after medical expenses are realized and ensures that households can finance at least $\underline{c}$ while respecting the borrowing constraint. In the benchmark, $\underline{c}$ is calibrated to empirical estimates of consumption floors used in the medical-expenditure and retirement literature, and sensitivity analysis varies this parameter because it governs the insurance value of public transfers.

Jointly calibrated parameters

The second step jointly calibrates parameters that govern preferences, labor supply, bequests, and the fiscal closure. The discount factor β targets the aggregate wealth-to-output ratio. The labor-disutility parameter θ_ℓ targets average hours or employment by age. The curvature parameters σ and χ are set to standard values or chosen to match consumption and labor-supply elasticities. The bequest parameters $(\theta_b, \kappa_b, \gamma_b)$ target the bequest-to-output ratio and the upper tail of the wealth distribution. Parameters governing retirement or claiming frictions are chosen to match claiming and retirement profiles by age and education.

3.4 Policy Experiments

The calibrated model is used to study pension reforms during demographic transition. The benchmark demographic transition changes fertility, survival, or both, thereby in-

Table 4: Calibration Strategy

Object	Parameter / process	Target or source
Panel A: First step — estimated outside the model or set by institutional rules		
Education shares	π_c, π_n	Cohort education distribution
Health transitions	$\Pi_j^h(h' \mid h, s)$	PSID, by age and education
Survival	$\psi_j(s, h)$	PSID, aligned with life tables
Medical spending	$\Pi_j^m(m \mid s, h)$	MEPS out-of-pocket expenditure distribution
Age efficiency	η_j^c, η_j^n	PSID earnings profiles by education
Productivity risk	$\Pi_j^z(z' \mid z, s)$	PSID panel earnings dynamics
Health productivity	ω_g, ω_b	Earnings gap by health status
Social Security	$\tau^{ss}, \bar{y}^{ss}, b(e, j^R)$	Statutory rules
Income taxes	$T(y)$ parameters	Average and marginal tax rates
Consumption floor	$\underline{c}$	Safety-net and medical-spending literature
Panel B: Second step — jointly calibrated in equilibrium		
Discounting	β	Wealth-to-output ratio
Labor disutility	θ_ℓ, χ	Hours/employment by age
Bequests	$\theta_b, \kappa_b, \gamma_b$	Bequests and wealth concentration
Claiming/retirement frictions	Taste-shock parameters	Claiming and retirement profiles by age

Notes: Panel A parameters are estimated directly from micro data (PSID, MEPS) or set from statutory rules, as described in Section 3. Panel B parameters are calibrated jointly so that the stationary equilibrium matches the listed moments.

creasing the old-age dependency ratio and generating pressure on the pay-as-you-go pension budget. We consider reforms that restore fiscal balance using alternative instruments: payroll-tax increases, benefit reductions, increases in the early or normal claiming ages, and combinations of these policies. For each reform, we compute the transition path of prices, quantities, government budgets, distributions, and welfare.

4 Results

This section reports results from a preliminary version of the model that abstracts from the health and education heterogeneity described in Section 2. The economy is calibrated to the United States, and the demographic shock is an increase in survival probabilities that raises life expectancy at age 18 from roughly 82 to 87 years and the old-age dependency ratio from 0.28 to 0.34. We compare four reforms against a common baseline in which the aging shock occurs and the consumption tax τ_c is the residual instrument that closes the government budget: (i) raising the payroll tax rate, (ii) cutting benefits proportionally, (iii) raising the NRA from 67 to 69 with cohort-based grandfathering, and (iv) removing the cap on payroll contributions while leaving the AIME cap unchanged.

4.1 Steady States

Table 5 compares the initial steady state with the post-aging steady states under each reform. Aging raises capital by 14 percent and lowers the interest rate by roughly 80 basis points, as longer expected lifetimes induce additional precautionary and life-cycle saving. Absent reform, closing the budget requires the consumption tax to rise from 6.4 to 8.8 percent. The three reforms that operate directly on the pension system—raising the payroll tax, cutting benefits, and raising the NRA—each return τ_c to approximately its pre-aging level. Uncapping contributions raises too little revenue to do so, leaving τ_c at 8.8 percent.

Cutting benefits and raising the NRA generate nearly identical long-run aggregates. This is the first indication that the NRA reform works primarily as an implicit benefit cut through the PIA adjustment rather than through later retirement: the average retirement age rises only from 64.0 to 64.1 years, consistent with the small retirement elasticity we estimate from the HRS. Because the retirement margin barely moves, the fiscal savings come almost entirely from lower benefits per retiree.

4.2 Transition Welfare

We measure welfare along the full general equilibrium transition. For each household alive at the date of reform, the consumption-equivalent variation λ is the permanent percentage change in consumption that equates lifetime utility under the reform to lifetime utility under the consumption-tax baseline, so that $\lambda > 0$ means the reform is preferred to raising the consumption tax. Table 6 reports averages weighted by the initial distribution.

Table 5: Steady State Comparison

	Initial SS	Aging $+\tau_c$	Raise tax	Cut benefits	Raise NRA	Uncap SS
Interest rate r	7.01%	6.22%	6.46%	6.00%	5.97%	6.24%
Wage w	3.04	3.16	3.12	3.19	3.20	3.15
Capital K	3.00	3.43	3.32	3.53	3.55	3.41
Output Y	1.00	1.07	1.06	1.08	1.08	1.07
Consumption C	0.75	0.80	0.79	0.80	0.80	0.79
Effective labor N	0.21	0.22	0.22	0.22	0.22	0.22
Aggregate benefits	0.10	0.13	0.13	0.11	0.11	0.13
Consumption tax τ_c	6.4%	8.8%	6.4%	6.4%	6.3%	8.8%
Dependency ratio	0.28	0.34	0.34	0.34	0.34	0.34
Average retirement age	64.0	64.0	64.0	64.0	64.1	64.0

Notes: Output in the initial steady state is normalized to one. The "Aging $+\tau_c$" column is the reform-free baseline in which the consumption tax is the residual fiscal instrument. All other columns hold τ_c at its pre-aging level except "Uncap SS," where τ_c closes the remaining budget residual.

Raising the payroll tax is the only reform that dominates the baseline, with an aggregate CEV of +0.38 percent and support from 82 percent of the existing population. Both workers (+0.33 percent) and retirees (+0.52 percent) gain: the payroll tax distortions are smaller than those of the consumption tax it substitutes for, and retirees gain the most because they pay the consumption tax on their benefits but no payroll tax. Cutting benefits is the costliest reform (−0.61 percent) and its burden is sharply uneven: retirees lose 1.24 percent against 0.37 percent for workers, only the youngest workers—who avoid decades of higher consumption taxes before retiring—gain, and support falls to 25 percent.

Raising the NRA (−0.18 percent) shifts the burden onto workers (−0.29 percent) while leaving retirees marginally better off (+0.11 percent). The age profile is informative: mid-career workers aged 36–55 bear the largest cost (−0.45 percent), facing the full PIA reduction with limited time to adjust savings, whereas workers aged 56–61 are nearly indifferent under cohort-based grandfathering. Uncapping contributions is approximately welfare-neutral in the aggregate (−0.06 percent) but is supported by almost no one: it raises too little revenue to lower τ_c while imposing a concentrated cost on top earners who receive no offsetting benefit increase.

Decomposing these CEVs into a direct reform channel, a consumption-tax channel, and a general equilibrium price channel shows that the first two dominate; differences in equilibrium price paths across reforms contribute at most 0.07 percentage points. The welfare ranking is therefore driven by the incidence of the fiscal instruments themselves rather than by their differential effects on factor prices.

Taken together, the results point in three directions. The low retirement elasticity implies that raising the NRA restores solvency by reducing benefits rather than by inducing

Table 6: Transition Welfare (CEV, existing cohorts)

Panel A: Aggregate and by labor force status				
Reform	Aggregate	Workers	Retirees	Support
Raise tax	+0.38%	+0.33%	+0.52%	81.6%
Cut benefits	−0.61%	−0.37%	−1.24%	25.1%
Raise NRA	−0.18%	−0.29%	+0.11%	35.6%
Uncap SS	−0.06%	−0.08%	−0.03%	0.7%

Panel B: By age group					
Reform	18–35	36–55	56–61	62–69	70+
Raise tax	−0.06%	+0.51%	+0.73%	+0.68%	+0.44%
Cut benefits	+0.38%	−0.63%	−1.25%	−1.37%	−1.19%
Raise NRA	−0.25%	−0.45%	−0.05%	+0.16%	+0.10%
Uncap SS	−0.12%	−0.06%	−0.03%	−0.03%	−0.04%

Notes: CEV is measured relative to the baseline in which the aging shock occurs and the consumption tax adjusts to close the government budget. Positive values indicate that the reform is preferred to consumption-tax adjustment. Support is the mass of the existing population with CEV $>$ 0, computed at the bin level over age, wealth, productivity, and AIME and weighted by the initial steady state distribution. Households aged 62–69 face the retirement choice.

later retirement, and should be evaluated as such. The reforms differ far more in who bears the burden than in the aggregates they deliver: cutting benefits and raising the NRA yield nearly identical steady states but very different welfare profiles across ages. And no single instrument is uniformly preferred once incidence across cohorts is accounted for, which motivates the Ramsey problem in Section 5: the welfare-maximizing reform is likely to combine instruments and to vary over the transition.

5 The Ramsey Framework

To determine the optimal response to demographic transition, we consider the problem of a benevolent Ramsey planner who chooses a sequence of policy instruments to maximize social welfare. The planner takes as given the competitive-equilibrium behavior described in Section 2.8 and chooses policy paths that affect household budget constraints, benefit rules, and government budgets.

Let Θ_t denote the vector of policy instruments in period t. In applications, Θ_t may include the payroll tax τ_t^{ss}, parameters of the benefit formula $b_t(\cdot)$, the early and normal claiming ages, income-tax parameters in $T_t(\cdot)$, government debt D_{t+1}, or the consumption floor $\underline{c}_t$. For any candidate policy sequence $\Theta \equiv \{\Theta_t\}_{t=0}^{\infty}$, the associated competitive equilibrium delivers household value functions, distributions, prices, aggregates, and gov-

ernment budgets.

The planner's objective can be written as a weighted sum of the welfare of households alive at the time of reform and future newborns:

$$\mathcal{W}(\Theta) = \sum_{j=1}^{J} \sum_{s\in\{c,n\}} \int \lambda_{j,s,0}(x) V_{j,0}^{\Theta}(x,s)\, d\mu_{j,s,0}(x) + \sum_{t=1}^{\infty} \beta_g^t \sum_{s\in\{c,n\}} \pi_{s,t} \int V_{1,t}^{\Theta}(x,s)\, d\mu_{1,s,t}(x), \tag{25}$$

where x collects the non-education individual states, $\lambda_{j,s,0}(x)$ are Pareto weights for households initially alive, β_g is the planner's discount factor over future cohorts, and $\mu_{1,s,t}$ is the initial distribution of newborns of type s at date t. A special case is an ex-ante newborn welfare criterion, which assigns welfare weight only to future newborns before their permanent education type and idiosyncratic risks are realized.

The Ramsey problem is

$$\max_{\Theta} \ \mathcal{W}(\Theta) \tag{26}$$

subject to household optimization, firm optimization, market clearing, the law of motion for the distribution, and the government and pension-system budget constraints. Computationally, each candidate policy sequence requires solving the full transition path: household problems are solved backward from a terminal steady state, distributions are simulated forward, and prices and fiscal variables are iterated until all equilibrium conditions hold.

6 Conclusion

We study how the burden of adjusting a pay-as-you-go pension system to population aging should be allocated across generations, combining an overlapping generations model with incomplete markets and endogenous retirement, disciplined by retirement elasticities estimated from the HRS. The estimated elasticity is small, and this shapes the evaluation of every reform we consider: raising the normal retirement age restores solvency not by inducing later retirement, but by reducing benefits through the PIA adjustment. In our transition experiments, raising the payroll tax is the only reform preferred to closing the budget with a consumption tax, and it is supported by a large majority of households alive at the time of reform. Cutting benefits delivers aggregates nearly identical to raising the NRA but concentrates the cost on current retirees. The reforms thus differ far more in their incidence than in their macroeconomic consequences.

Our agenda from here has two parts. The first is to bring the richer heterogeneity developed in Section 2 into the quantitative experiments. Because health and longevity are correlated with lifetime earnings, adding education types, stochastic health, and medical-expenditure risk should sharpen the distributional consequences of benefit cuts and NRA increases in particular. The second is to move beyond the single-instrument reforms studied here, none of which is uniformly preferred, to the Ramsey problem of Section 5:

we solve for the welfare-maximizing, time-varying combination of instruments and ask whether the resulting reform path can command majority support.

References

Attanasio, O., S. Kitao, and G. L. Violante (2007). Global demographic trends and social security reform. *Journal of Monetary Economics 54*(1), 144–198.

Auerbach, A. J. and L. J. Kotlikoff (1987). *Dynamic Fiscal Policy*. Cambridge: Cambridge University Press.

Bairoliya, N. and K. McKiernan (2021). Revisiting retirement and social security claiming decisions. *Available at SSRN 3896031*.

Conesa, J. C. and D. Krueger (1999). Social security reform with heterogeneous agents. *Review of Economic Dynamics 2*(4), 757–795.

De Nardi, M., S. İmrohoroğlu, and T. J. Sargent (1999). Projected us demographics and social security. *Review of Economic Dynamics 2*(3), 575–615.

Dyrda, S. and M. Pedroni (2023). Optimal fiscal policy in a model with uninsurable idiosyncratic income risk. *The Review of Economic Studies 90*(2), 744–780.

Díaz-Giménez, J. and J. Díaz-Saavedra (2009). Delaying retirement in spain. *Review of Economic Dynamics 12*(1), 147–167.

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.

Hosseini, R. and A. Shourideh (2019). Retirement financing: An optimal reform approach. *Econometrica 87*(4), 1205–1265.

Huang, H., S. İmrohoroğlu, and T. J. Sargent (1997). Two computations to fund social security. *Macroeconomic Dynamics 1*(1), 7–44.

Huggett, M. and J. C. Parra (2010). How well does the us social insurance system provide social insurance? *Journal of Political Economy 118*(1), 76–112.

Imrohoroglu, A., S. İmrohoroğlu, and D. H. Joines (1995). A life cycle analysis of social security. *Economic Theory 6*(1), 83–114.

İmrohoroğlu, S. and S. Kitao (2012). Social security reforms: Benefit claiming, labor force participation, and long-run sustainability. *American Economic Journal: Macroeconomics 4*(3), 96–127.

Jones, J. B. and Y. Li (2023). Social security reform with heterogeneous mortality. *Review of Economic Dynamics 48*, 320–344.

Kitao, S. (2014). Sustainable social security: Four options. *Review of Economic Dynamics 17*(4), 756–779.

Krueger, D. and F. Kubler (2006). Pareto-improving social security reform when financial markets are incomplete!? *American Economic Review 96*(3), 737–755.

Makarski, K., J. Tyrowicz, and O. Komada (2024). Capital income taxation and reforming social security in an olg economy. *Journal of Economic Dynamics and Control 165*, 104878.

McGrattan, E. R. and E. C. Prescott (2017). On financing retirement with an aging population. *Quantitative Economics 8*(1), 75–115.

Ndiaye, A. and Z. Yu (2025). Flexible retirement and optimal taxation. Technical report, CESifo Working Paper.

Nishiyama, S. and K. Smetters (2007). Does social security privatization produce efficiency gains? *The Quarterly Journal of Economics 122*(4), 1677–1719.

Rogerson, R. and J. Wallenius (2013). Nonconvexities, retirement, and the elasticity of labor supply. *American Economic Review 103*(4), 1445–1462.

Shourideh, A. and M. Troshkin (2017). Incentives and efficiency of pension systems. *Unpublished Discussion Paper*.

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