Social Security and Unfunded Government Debt

Elliot Spears[†]

January 2026

Abstract

I study how population aging affects macroeconomic outcomes when rising Social Security obligations are financed under different fiscal-monetary regimes. I build a life-cycle model with workers and retirees, capital accumulation, pay-as-you-go transfers, government debt, and a fiscal-valuation channel through which unfunded liabilities affect inflation and the real value of debt. Using current U.S. demographic projections, I simulate a medium-run retirement wave and compare funded versus unfunded financing of the induced transfer burden. The results show that financing assumptions strongly shape the adjustment: without reform to the pension program, unfunded accommodation produces the largest inflation response. Benefit cuts under a fully-funded regime reduce fiscal pressure, but redistribute welfare significantly across retirees, workers, and future generations. Immediate and gradual cuts paired with capital-income-tax adjustments rank first and second, respectively, and prove to be robust to sensitivity analyses. No single benefit-reform package benefits every identified group: grandfathering protects incumbent retirees by shifting losses toward near-term entrants, while unfunded reforms benefit current generations at future cohorts' expense.

[†] University of Washington, Department of Economics. Email: espear1@uw.edu.

1 Introduction

This paper analyzes how future social security expenditures in the United States will affect the economy under various approaches to financing the outlays. Specifically, I assume that a fiscal expansion will be necessary in order to finance the future retirement wave and then compare the effects of unfunded versus funded fiscal expansions. The distinction between "unfunded" and "funded" fiscal expansions has to do with how the government responds to an increase in transfers. In the unfunded case, the government does not make the necessary budgetary adjustments (say, via taxes) to cover the added spending. In the fully funded case, they do. The consequence of this is that, in the unfunded case, there is no current or future revenue stream in place that will cover any debt accumulated to finance the transfers, which leads to movements in the price level to reestablish a balance between the real value of nominal debt and the present value of future expected surpluses Cochrane (2023)

To this end, I combine ingredients from existing models in order to create an overlapping generations model with nominal rigidities and demographic dynamics. The fiscal structure that allows me to separate funded and unfunded regimes comes from Bianchi et al. (2023), where the funded case corresponds to what is commonly known as an active-monetary/passive-fiscal regime, where the government is expected to balance its budget while the monetary authority controls inflation. Unfunded transfers correspond to the active-fiscal/passive-monetary case where the central bank accommodates unfunded fiscal expansions. The model in this paper preserves the same "monetary-led" and "fiscally-led" policy mix that appears in Bianchi et al. (2023).

The overlapping generations (OLG) ingredients that give us the framework necessary to introduce social security and retirement dynamics comes from the canonical Gertler (1999), where there are two kinds of agents: retirees and workers. Throughout this paper, I assume the special case of the Gertler OLG model where labor is supplied elastically. To simulate retirement waves of the kind the U.S. faces, I import the demographic modeling features similar to Kara and Von Thadden (2010). Their paper is calibrated and tailored to match moments from the Eurozone. This paper, of course, is calibrated for the U.S.

By incorporating monopolistically competitive firms into the model, I am able to generate nominal rigidities and graft these ingredients into a New Keynesian (NK) framework. The nominal rigidities are imported into the OLG setup in a manner that follows Fujiwara and Teranishi (2008).

In the model, the retirees are allowed to work (albeit with lower productivity than workers), yet only workers are subject to a Social Security payroll tax on labor that they supply elastically. The workers' payroll taxes finance all benefits of the PAYGO program in the fully-funded case. In the unfunded scenario, the payroll tax does not adjust to meet the shortfall in revenue necessary to finance the fiscal expansion. The size of the benefits follows Kara and Von Thadden (2010) in that they are tied to wages via the replacement rate (and individual's Social Security benefits relative to preretirement earnings).

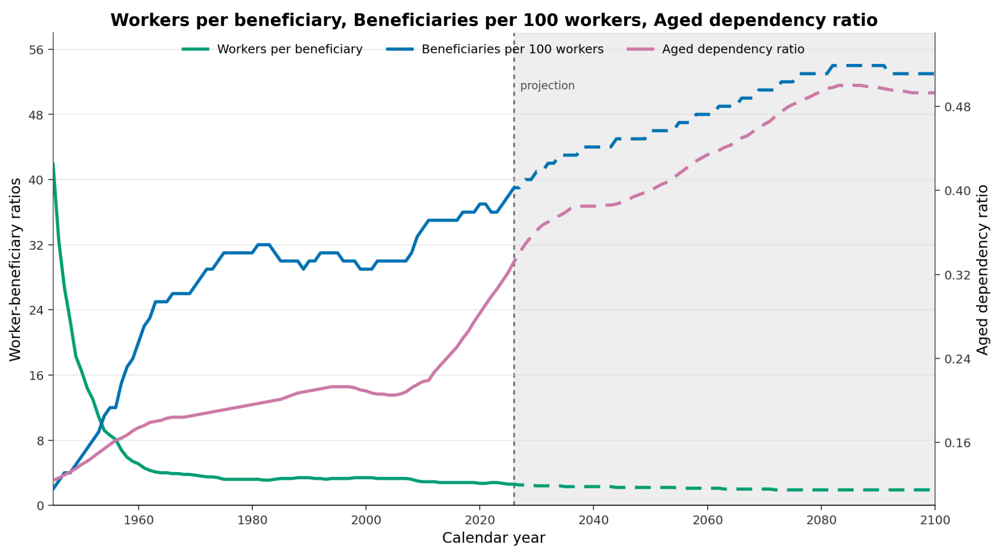

Figure 1: Worker-Beneficiary and Aged Dependency Ratios

Source: SSA 2026 OASDI Trustees Report, Tables IV.B4 and V.A3

The real-world phenomenon of interest that is central to this paper is the projected retirement wave of the last remnants of the baby-boom generation. According to the Social Security Administration (SSA), the aged dependency ratio was 0.322 as of 2025, and is projected to increase to 0.379 by 2035. [1] This means the retiree share of the U.S. population

[1]The aged dependency ratio is the ratio of the population aged 65 and over to the population aged 20–64. See Table V.A3 of the 2026 OASDI Trustees Report.

is projected to grow from 24.4% as of 2025 to 27.5% by 2035,[2] a 3.1% increase (12.7% relative to 2025). After calibrating the model to match moments in the U.S. as of 2025, I simulate this 10-year retirement wave and analyze the macroeconomic effects of financing the necessary federal outlays via unfunded debt, versus a fully funded arrangement via payroll taxes. Figure 1 provides a visualization of the projections for the aged dependency ratio (magenta) and also provides the series for the ratio of workers per Social Security beneficiary (green) as well as the beneficiaries per 100 workers (blue).

In the funded case, taxes rise substantially, which leaves the ratio of primary surpluses-to-output relatively constant over all horizons up to 60 years. Due to the fact that the retiree-share rises, and retirees have a larger marginal propensity to consume (MPC) relative to workers, the rise in retirees and corresponding expansion to social security raises overall consumption. Worker consumption deviates from steady state, with a trough of about -1.75% 12 years into the future, and never fully recovers. The rise in retiree consumption more than offsets the fall in that of workers, for a net rise in aggregate consumption. The resulting wealth effects of the fall in consumption for workers dominates the substitution effects of the rise in the distortionary payroll tax in determining workers' labor supply. On net, workers raise their labor supply indefinitely, while retiree labor supply falls below steady state for a transient period.

An unfunded expansion of the pension system, by contrast, results in a meaningful and persistent rise in inflation (nearly 2% on impact and remaining above trend for the entire horizon). Aggregate consumption falls immediately in this case by about 1.8% on impact, before rebounding and even rising above the fully funded benchmark consumption level after 10 years, which is entirely owing to the fact that worker consumption recovers quickly after its initially large fall in the immediate aftermath of the shock. The primary surplus-to-output ratio declines and never recovers, hence necessitating the inflationary adjustments.

I build upon these baseline results to run several policy counterfactuals. I compare the baseline unfunded case (assuming no reform to Social Security), to cases involving immediate across-the-board cuts to the program, gradual cuts, and future-retiree-only cuts. I also mix

[2] The retiree share is simply aged dependency ratio / (1+ aged dependency ratio) $\equiv \frac{POP_{65+}/POP_{20-64}}{1+(POP_{65+}/POP_{20-64})}$ $= \frac{POP_{65+}}{(POP_{20-64}+POP_{65+})}$, where "$POP_{65+}$" is the population aged 65 and over.

in these cuts with varying degrees of tax hikes that leave varying proportions of the fiscal expansion unfunded (allowing for some inflation) to assess what the optimal path forward for the program is. Following previous work, I also introduce a maturity structure to the debt to allow for long-term debt, which is known to alter the dynamics of inflation and other macro aggregates Cochrane (2001). I find that, in general, a commitment to fully fund the debt dominates unfunded fiscal adjustments in terms of welfare effects. I also find that an immediate cut to benefits, accompanied by a reduction in capital tax rates made possible by the fiscal space that the benefit cut provides, is the best path forward among the all reforms assessed under both fiscal regimes (funded and unfunded).

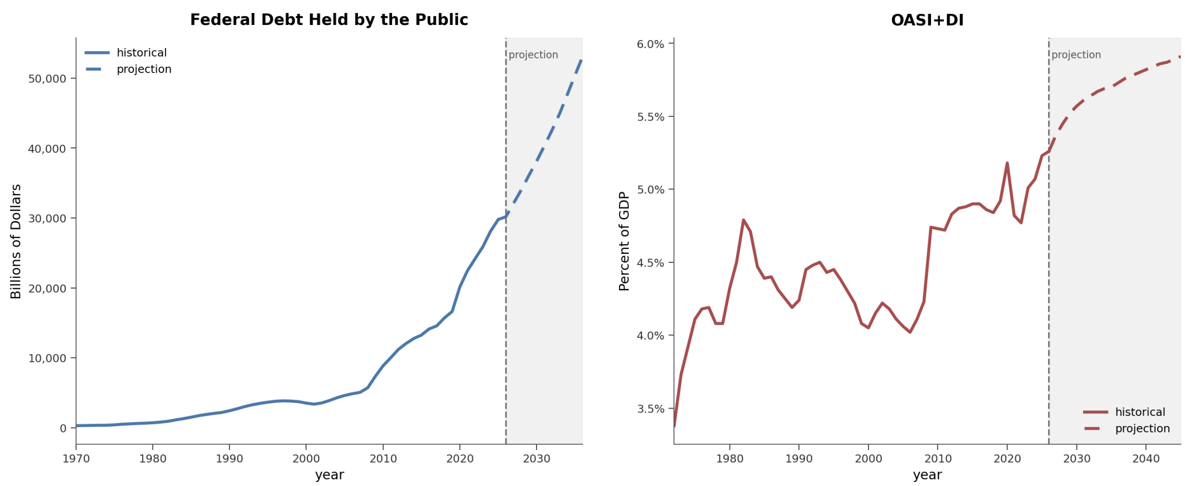

Figure 2: Debt-GDP and Social Security Projections

The sources for this plot are the CBO and SSA.

This paper is written at a time of increased fiscal strain owing to the deficits incurred during the COVID-19 pandemic, on top of the rising costs of entitlement programs in the U.S. (such as Social Security and Medicare/Medicaid). Figure 2 provides an illustration of the problem. The left panel of the figure shows the amount of federal debt held by the public in billions of dollars. The 2008 financial crisis demarcates a regime change that is subsequently followed by the massive fiscal expansion (owing to the ARPA and CARES Act) in the wake of the pandemic. The vertical dashed bar separates the end of historical data and the beginning of projections by the Congressional Budget Office (CBO), which represents the gray shaded region. The right panel represents the historical and projected outlook (according to the

SSA) of the Old-Age & Survivors Insurance (OASI) and Disability Insurance (DI) programs, more commonly discussed jointly as (OASDI). The impending retirement wave is projected to take up one additional percent of GDP over a 20 year horizon as retirement rises steadily over the the next 15 years.

Consequently, this paper is situated at the intersection of several large, but increasingly relevant strands of literature. The first literature relates to monetary-fiscal linkages, especially fiscal dominance and its influence on the price level Sargent (1981); Leeper (1991); Sims (1994); Schmitt-Grohé and Uribe (2000); Cochrane (2001); Sims (2011); Bianchi et al. (2023); Cochrane (2023). A second literature of relevance is the public finance literature addressing issues related to Social Security in the U.S. The foundational text for this literature, and the paper at hand, is Gertler (1999), which constructs an OLG model with workers and retirees, where the former faces some probability of retirement, the latter some probability of death and, in the meantime, receives pension benefits. Gertler's work builds upon that of Blanchard (1985); Romer (1988); Weil (1989), where the so-called "Blanchard-Weil" framework enables one to examine the effects of wealth re-distributions across generational cohorts. The problem is that this framework omits the kind of life-cycle framework necessary to capture the net-wealth effects of government debt Barro (1974). Extensions of Gertler have been manifold, with Carvalho et al. (2016) introducing demographic features such as population growth and changes in longevity. Work on optimal pension reform has also been burgeoning, with work by McGrattan and Prescott (2017); Hosseini and Shourideh (2019) examining what fiscal adjustments would improve the program. The latter paper uses Pareto optimal policy reforms to conduct an analysis leading to the identification of a lack of asset subsidies for retirees as a source of inefficiency in the current retirement financing arrangement.

2 The Model

The model involves a blend of the overlapping generations and life-cycle framework for households from Gertler (1999) and the fiscal-monetary mix on the government side from Bianchi et al. (2023). Agents enter life as workers who supply their labor elastically. Their wages are subject to a distortionary payroll tax that finances the pension system for retirees (which workers have a set probability of becoming each period). Retirees may also work, though they're less productive and their labor is not subject to the payroll tax. Any additional taxes are lump-sum and are extracted from the workers.

Workers own a monopolistically competitive firm which produces its good using capital and labor and receive capital income through a mutual fund, which also holds government debt. The government purchases goods and makes transfers (in the form of pension benefits). It also issues debt. When the government embarks on an unfunded fiscal expansion (deficit spending without corresponding future surpluses), inflation adjusts so as to revaluate outstanding government debt.

2.1 Households

Individuals are born into the economy as a workers. Those who are workers as of date t remain a workers going into $t+1$ with probability ω_t. They face potential retirement with probability $1-\omega_t$. Once someone enters into retirement, they survive into the next period with probability γ and die with probability $1-\gamma$. The basic demographic structure is

$$N^w_{t+1} = (1-\omega_t+n)N_t + \omega_t N_t = (1+n_t)N^w_t \tag{1}$$

$$N^r_{t+1} = (1-\omega_t)N^w_t + \gamma N^r_t \tag{2}$$

Equation (1) represents the evolution of the workforce, while equation (2) is the law of motion for retirees.

Let

$$\psi_t \equiv \frac{N^r_t}{N^w_t} \tag{3}$$

denote the retiree–worker ratio. Equations (1) and (2) imply

$$\psi_{t+1} = \frac{1-\omega_t+\gamma\psi_t}{1+n_t} \tag{4}$$

To simulate a retirement-wave shock, I use the the probability of remaining a worker (ω_t) as the variable of interest for this purpose. If we have some retirement process d_t, we have:

$$d_t = \rho_d d_{t-1} + \varepsilon_t^R \tag{5}$$

$$\omega_t = \bar{\omega} - d_t \tag{6}$$

The experiment holds both individual longevity as well as the worker population growth rate fixed $\gamma = \bar{\gamma},\ n_t = \bar{n}$. A positive shock through ε_t^R thus accelerates transitions into retirement, thereby raising ψ_{t+1}.

Let $\ell_t^j \in (0,1)$ denote hours and $1 - \ell_t^j$ leisure for an individual of type $j \in \{w, r\}$. Individual preferences follow Gertler's CES nonexpected utility form with an elastic labor supply.:

$$V_t^r = \left[((C_t^r)^\nu (1-\ell_t^r)^{1-\nu})^\rho + \beta\gamma (V_{t+1}^r)^\rho\right]^{1/\rho} \tag{7}$$

$$V_t^w = \left[((C_t^w)^\nu (1-\ell_t^w)^{1-\nu})^\rho + \beta\left(\omega_t V_{t+1}^w + (1-\omega_t) V_{t+1}^r\right)^\rho\right]^{1/\rho} \tag{8}$$

With intertemporal elasticity of substitution parameter $\sigma = \frac{1}{1-\rho}$, where parameter ρ determines the degree of consumption smoothing. These preferences make agents risk neutral with respect to income risk, which helps offset the effects of an exogenously determined retirement probability.

The model does not insure against the income loss at retirement because this would eliminate the exact kind of life-cycle spending-savings behavior that I am trying to model. There is, however, a kind of death insurance in the form of an annuity market that invests retirees' savings. The wealth of retirees who die each period is disbursed across the surviving retirees, so a surviving retiree earns a return R_t^A/γ_{t-1}.

Workers and retirees face different budget constraints. The worker faces a payroll tax, which the retiree is not subject to. Profits from the firm are solely alloted to the worker. The worker also faces the prospect of lump-sum taxes, which the retiree is exempt from. A_t^j represents previously acquired assets carried into period t. The worker's constraint is represented by:

$$C_t^w + A_{t+1}^w = R_t^A A_t^w + (1-\tau_t^p) W_t^w \ell_t^w + DIV_t - T_t \tag{9}$$

where τ_t^p is the payroll-tax rate, T_t is a residual lump-sum tax that covers anything not covered by the payroll tax, and DIV_t denotes any dividend profits from the firm. In the implemented competitive real block profits are zero and residual taxes are assigned to workers.

The retiree's assets evolve according to:

$$C_t^r + A_{t+1}^r = \frac{R_t^A}{\gamma} \; A_t^r + \xi W_t^r \ell_t^r + E_t^r \tag{10}$$

where $0 < \xi < 1$ is retiree labor productivity and E_t^r is the individual pension benefit.

The first-order conditions from (7)–(10) give us:

$$\ell_t^j = 1 - \varsigma \frac{c_t^j}{w_t^j} \qquad \text{for} \qquad j \in \{w, r\} \tag{11}$$

Where $w_t^w \equiv (1 - \tau_t^p) w_t$ and $w_t^r \equiv \xi w_t$. Along with the following consumption Euler equation for individual retirees:

$$\frac{C_{i,t+1}^r}{C_{it}^r} = \left(\frac{W_t^r}{W_{t+1}^r} \right)^{(1-\nu)\rho\sigma} \beta^\sigma \left(R_{t+1}^A \right)^\sigma \tag{12}$$

and workers:

$$\frac{C_{i,t+1}^w}{C_{it}^w} = \left(\frac{W_t^w}{W_{t+1}^w} \right)^{(1-\nu)\rho\sigma} \beta^\sigma \left(R_{t+1}^A \right)^\sigma \Omega_{t+1}^{\sigma-1} \tag{13}$$

The term Ω_{t+1} is a weight that adjusts endogenously to the likelihood of retirement, MPC, and an agent's productivity while in retirement: $\Omega_t = \omega_t + (1 - \omega_t) \chi_t \epsilon_t^{1/(1-\sigma)}$. In terms of resource allocation, this term affects intertemporal consumption decisions for workers by forcing them to take into account the imminence of retirement, which gives rise to a savings motive for workers in preparation for their transition out of the workforce.

The next step is to conjecture solutions to the consumption functions for workers and retirees. For the retiree we have: $C_{it}^r = \epsilon_t \mu_t \left[\frac{R_t^A}{\gamma_{t-1}} A_{it}^r + H_{it}^r + S_{it} \right]$, and for workers we get: $C_{it}^w = \mu_t \left[R_t^A A_{it}^w + H_{it}^w + S_{it}^w \right]$. The term μ_t is the MPC for workers, while $\epsilon_t \mu_t$ is the MPC for retirees. Since retirees have higher MPCs than workers, $\epsilon_t > 1$.

Call λ_t the share of financial wealth held by retirees (excluding social security payments). H_{it}^j represents the present discounted value of all future labor income for individual j. In the aggregate, this becomes

$$H_t = \sum_{v=0}^{\infty} \frac{N_{t+v} W_{t+v}}{\Pi_{z=1}^{v}(1+n) R_{t+z}^{A} \Omega_{t+z}/\omega_{t+z}}$$

S_{it} is the per-retiree i value of social security benefits paid out at time t. Each worker i factors into their consumption decisions the expected social security benefits they anticipate receiving upon retirement, S_{it}^{w}. If we define $\hat{S}_t = \frac{S_t}{\psi_t N_t}$ as the per-retiree social security payment at time t, then we end up with an expression for S_t^w that represents the total value of the social security benefits for all laborers at time t:

$$S_t^w = \sum_{v=0}^{\infty} \frac{[(1-\omega_{t+v})\Pi_{j=0}^{v-1}\omega_{t+j}] N_t \left(\frac{\epsilon_{t+1+v}\hat{S}_{t+1+v}}{R_{t+v}^{A}\Omega_{t+v}}\right)}{\Pi_{z=1}^{v} R_{t+v}^{A} \Omega_{t+v}}$$

Combining conditions, we end up with the following equations MPCs:

$$\mu_t = 1 - \left(\frac{w_t^w}{w_{t+1}^w}\right)^{(1-\nu)\rho\sigma} \beta^{\sigma} (R_{t+1}^{A}\Omega_{t+1})^{\sigma-1} \frac{\mu_t}{\mu_{t+1}} \tag{14}$$

$$\epsilon_t \mu_t = 1 - \left(\frac{w_t^r}{w_{t+1}^r}\right)^{(1-\nu)\rho\sigma} \beta^{\sigma} \gamma (R_{t+1}^{A})^{\sigma-1} \frac{\epsilon_t \mu_t}{\epsilon_{t+1}\mu_{t+1}} \tag{15}$$

To derive the law of motion for retiree-share of financial wealth, I specify aggregate retiree saving before workers enter into retirement as follows:

$$SAV_t^r = \lambda_t R_t^A A_t + (1-\tau^p)\xi w_t \psi_t \ell_t^r + E_t - C_t^r \tag{16}$$

Where E_t is the aggregate value of benefits for current retirees. A similar end-of-period saving exist for workers:

$$SAV_t^w = (1-\lambda_t) R_t^A A_t + (1-\tau_t^p) w_t \ell_t^w - \bar{T} - C_t^w \tag{17}$$

With lump-sum taxes $\bar{T}$. To derive the law of motion for the share of assets held by retirees (λ_T), note that the savings of existing retirees enter this equation in addition to the fraction $(1-\omega_t)$ of workers that retire in a given period. Therefore, we have:

$$\lambda_{t+1}A_{t+1} = \lambda_t R_t^A A_t + (1-\tau_t^p)\xi w_t \psi_t \ell_t^r + E_t - C_t^r + (1-\omega_t)[(1-\lambda_t)R_t^A A_t + (1-\tau_t^p)w_t \ell_t^w - \bar{T} - C_t^w] \tag{18}$$

Given $\frac{A_t^r}{A_t}$. This is analogous to Gertler's wealth-distribution equation, but with retiree labor income. The model, up to this point, follows Gertler (1999) closely, with a few variations. One important difference is how R_t^A is determined. In the original Gertler framework, a single rate of interest prevailed that was determined endogenously by the productivity of capital, with no government influence. Here, there is a government and a central bank. The central bank sets a policy rate that affects asset returns. More on this below. Now I turn to the production side of the economy.

3 Firms

3.1 Final goods

Intermediate goods firms have unit mass and supply differentiated products, $z \in [0,1]$, that are transformed into a final good according to:

$$Y_t = \left[\int_0^1 y_t(z)^{\frac{(\varepsilon_p-1)}{\varepsilon_p}} dz\right]^{\frac{\varepsilon_p}{(\varepsilon_p-1)}}, \qquad \varepsilon_p > 1 \tag{19}$$

Where ε_p is the price elasticity of demand. The demand for differentiated product z is represented by:

$$y_t(z) = \left(\frac{P_t(z)}{P_t}\right)^{-\varepsilon_p} Y_t \tag{20}$$

with $P_t(z)$ representing the price of the intermediate good z and P_t denoting the price level. The aggregate price level follows the standard CES structure and is specified as:

$$\tag{21}$$

$$P_t = \left[\int_0^1 P_t(z)^{1-\varepsilon_p} dz \right]^{\frac{1}{(1-\varepsilon_p)}} \tag{22}$$

Define $H_t(i)$ as the amount of effective labor hours employed by intermediate-good firm i and let $K_t^s(i)$ be capital utilized in their production process. The aggregate specifications for these variables are:

$$H_t = N_t^w \ell_t^w + \xi N_t^r \ell_t^r = N_t^w L_t, \tag{23}$$

$$K_t^s = v_t K_t, \tag{24}$$

where v_t represents the rate of capital utilization for leased capital K_t, which determines the cost of each unit of physical capital employed in the production process.

3.2 Intermediate goods

Intermediate-good firm i faces Cobb–Douglas technology:

$$y_t(z) = [K_t^s(i)]^\alpha [X_t H_t(i)]^{1-\alpha} \tag{25}$$

The cost-minimization problem of the firm yields a real marginal cost that is independent of firm i, meaning it is the same across intermediate-good firms:

$$mc_t = \left(\frac{r_t^k}{\alpha} \right)^\alpha \left(\frac{w_t}{1-\alpha} \right)^{1-\alpha}. \tag{26}$$

The intermediate-good firms rent $v_t k_t$ units of effective capital each period at a cost per unit of physical capital of $\Phi(v_t)$. Capital income is taxed at a rate of $\tau_{K,t}$. Therefore, the intermediate-good firm chooses a utilization rate according to:

$$\max_{v_t \geq 0} \left\{ (1-\tau_{K,t}) r_t^k v_t k_t - \Phi(v_t) k_t \right\} \tag{27}$$

Following the specified functional form of $\Phi(v_t)$ in Bianchi et al. (2023), we end up with the following FOC:

$$(1-\tau_{K,t})r_t^k = (1-\bar{\tau}_K)\bar{r}^k v_t^{\phi_u/1-\phi_u} \tag{28}$$

Firms also face an adjustment cost associated with their investments, which affects the accumulation of physical capital. The details are included in the appendix.

Each intermediate-good firm is able to reset its price with probability $1-\theta_p$. If a firm does not get the opportunity to change its price in a given period, it must keep its price from the previous period. When the firm does get to reset its price, it chooses P_t^* so as to maximize the value of future expected profits once the new price is locked in until the next opportunity comes along to adjust it:

$$\max_{P_t^*} \ \mathbb{E}_t \sum_{s=0}^{\infty} \theta_p^s \Xi_{t+s} \left[P_t^* Y_{t+s} - MC_{t+s} Y_{t+s}\right], \tag{29}$$

where Ξ_{t+s} is the stochastic discount factor of the firms' owners. In our model, all firms are held by the Gertler mutual fund. Thus, the discount factor from t to $t+1$ is $1/R_{t+1}$.

3.3 Government

3.3.1 Fiscal Policy

Following Bianchi et al. (2023), the government finances its expenditures G_t, transfers (in the form of pension benefits) E_t, and residual debt via the raising of taxes $\tau_{i,t}$ on wages W, capital income K, and the levying of lump-sum taxes T_t. The lump-sum taxes are only ever levied on the workers and only come into effect to satisfy any shortfall in the tax revenue stemming from wages and capital. The government's budget constraint is therefore specified as:

$$B_t + \tau_{K,t} R_{K,t} v_t K_t + \tau_{L,t} W_t L_t + T_t = R_t^b B_{t-1} + G_t + E_t \tag{30}$$

If we call $s_{b,t}$ the debt-to-GDP ratio of the economy at time t, then the individual tax rates $\tau_{i,t}$ adjust according to:

$$\tau_{K,t} = \rho_K \tau_{K,t-1} + (1 - \rho_K)(\gamma_K s_{b,t-1} - \gamma_K^F s_{b,t-1}^F) \tag{31}$$

$$\tau_{L,t} = \rho_L \tau_{L,t-1} + (1 - \rho_L)(\gamma_L s_{b,t-1} - \gamma_L^F s_{b,t-1}^F) \tag{32}$$

The distinction between $s_{b,t}$ and $s_{b,t}^F$ requires a brief clarification of what these components represent. The central innovation of the Bianchi et al. (2023) framework is to introduce a fiscal-monetary mix that does not require that fiscal policy is always active, which it is in works like Cochrane (2023). This is accomplished by splitting the model into two "economies": the actual and shadow economies. The results of our modeling experiments use the shadow economy as a fully-funded baseline with which we can generate the funded components of debt and inflation that the federal government and central bank will respond to. In the actual economy, we have partially-funded or even fully-unfunded debt which generate an unfunded component to debt and inflation, which neither the federal government, nor the central bank may fully respond to. Therefore, the variable $s_{b,t}$ represents the debt-to-GDP ratio in the actual economy, and $s_{b,t}^F$ represents the debt-to-GDP ratio in the shadow counterpart. The simultaneous and differential responses to a funded and unfunded component of inflation and debt give us the coexistence of the monetary and fiscally-led mixes, where fiscal-dominance is activated in the case of unfunded social security expansions, again, following Bianchi et al. (2023).

3.3.2 Monetary Policy

The central bank responds to inflation and the output gap. The policy rule differs from the conventional Taylor rule in that the central banks responds differently to two separate sources of inflation: fiscal and non-fiscal inflation. The former arises when the government fails to raise the necessary surpluses to cover its future debt obligations. The central bank responds in the usual manner to non-fiscal inflation (more than one-for-one), while it "accommodates" fiscal inflation (less than one-for-one response). Accommodating fiscal inflation gives rise to necessary adjustments in the price level that ensue once the government commits to a deficit spending path that isn't backed by future tax revenue. The policy rule therefore follows:

$$r_{n,t} = \rho_r r_{n,t-1} + (1-\rho_r)\left[\phi_M \pi_t + \phi_F \pi_t^F + \phi_y y_t\right] + \varepsilon_t^{rn}$$

Where π_t^F represents fiscal inflation, which the central bank accommodates ($\phi_F < 1$). Note that the central bank responds in the usual fashion all other kinds of inflation as prescribed in the literature Woodford (2011), which means $\phi_M > 1$.

3.3.3 Government Debt

Let P_t^b represent the price of a long-term bond B_t issued today that pays ρ_b^j across j periods in the future. This notation follows both Bianchi et al. (2023) and Cochrane (2001)[3]. The return on the portfolio of long-term government debt is specified as:

$$R_{t,t+1}^b = \frac{1 + \rho_b P_{t+1}^b}{P_t^b}$$

With the no-arbitrage condition that the policy rate, the return on government debt, and the return on the mutual fund portfolio relate according to $R_{n,t} = \mathbb{E}_t\left[R_{t,t+1}^b\right] = R_t^A$.

4 Calibration

Some of the key parameters of this model are inherited from Gertler (1999). However, given that his paper was written two and a half decades prior to this one, most of those inherited parameters need to be updated to current data. To this end, I calibrate the model to match current and projected data as of 2026 from the SSA and CBO. Additionally, the main phenomenon of interest is the retirement wave of the baby-boom generation. Over the course of about 35 years, I simulate, year over year, an increase in the retirees per worker and aggregate social security benefits in accordance with data from the SSA.

According to the SSA, about 199.9 million people in the U.S. are between the ages of 20—64, while about 66.4 million are 65+ as of 2026. The old age dependency ratio is therefore at about 0.332. In one decade, that ratio is expected to grow to 0.38. In reality,

[3]For details of the geometric maturity structure of long-term debt, see the appendix of Cochrane (2022).

the dependency ratio is projected to peak in 2085 at 0.50. However, I focus my attention on nearer-term phenomena by simulating a single wave up until 2036 in order to isolate the effects of various reforms more precisely. Therefore, I simulate the dependency ratio wave as a smoothly increasing figure until 2036 (consistent with the data), at which point it smoothly declines.

Life expectancy is taken to be exogenously fixed in the model. Mean life expectancy at the age of 65 is currently 19.65 years. Therefore, I calculate a quarterly survival probability for retirees of $\gamma = 0.987$. Worker retention ω_t is technically an endogenous parameter that evolves according to: $\bar{\omega} - \varepsilon_t^{\text{ret}}$, where $\varepsilon_t^{\text{ret}}$ is an exogenous demographic shifter. The parameter $\bar{\omega}$ represents the probability that an individual remains in the worker "state" in a given quarter, as opposed to entering into retirement. This parameter is calibrated so that the quarterly probability of retirement is about 0.434%, which yields an annual retirement probability of about 1.73%.[4] The retirement-wave shock enters the model through ω_t.

Table 1: Current and Projected U.S. Demographic & Fiscal Parameters

Description	Parameter	Calibration	Source
Initial dependency ratio	$\bar{\psi}_{2026}$	0.332	SSA, Table V.A3
Peak dependency ratio	$\bar{\psi}_{2036}$	0.383	SSA, Table V.A3
Population growth	n	0.0003565/qtr.	SSA, Table V.A3
Retiree survival probability	γ	0.987/qtr.	—
Worker retention	ω	0.996/qtr.	—
Gov. purchases / GDP	G/Y	0.17	BEA, NIPA Table 1.1.5
OASDI cost / GDP	E/Y	0.0526	SSA, Figure II.D4
Debt held by public / GDP	B/Y	1.01	CBO

Table 1 provides an overview of the calibrations of the demographic features of the model (all calibrated to 2026 data and projections as of that year, along with some relevant steady-state calibrations on the fiscal side of the model.

In Table 2, I provide the calibrations for the structural parameters in the model. Most of the parameters are derived from Gertler (1999) and Bianchi et al. (2023). The latter paper uses Bayesian techniques to approximate many of their parameters, which I import

[4] This is calculated following the demographic structure of Kara and Von Thadden (2010), where $\psi = (1-\omega)/(g_{N,t} - \gamma)$, where $g_{N,t} \equiv 1 + n_t$ is the annualized growth rate of the working age population, which is about 1.00035% per quarter.

here as the data for the relevant parameters is recent enough to them reliable for the present study. The Gertler parameters that required an update are in the previous table. Therefore, the remaining calibrations from the social security and household block are imported from Gertler. The remainder of the relevant parameters are in the appendix and include tables for steady state results and experiment targets.

As I alluded to above, in simulating the imminent retirement-wave shock, the SSA provides year-over-year dependency ratio projections that, while increasing, exhibit a kind of jaggedness once implemented in the model, which generate uneven and choppy impulse responses (IRFs) in the endogenous variables. To ameliorate this issue, I smooth out the retirement wave, while maintaining consistency with the projected path in the data.

Table 2: Model Calibration

Description	Parameter	Calibration
Discount Factor	β	1
Intertemporal Elasticity of Sub.	σ	0.25
Consumption Weight	ν	0.40
Labor Efficiency of Retirees	ξ	0.60
Quarterly Worker Retention	$\bar{\omega}$	0.99566
Capital Share in Prod.	α	0.333
Depreciation Rate	δ	0.026
Price Markup	μ_p	1.14
Elasticity of Sub. (Goods)	ε_p	8.14
Calvo Price Non-Reset Prob.	θ_p	0.867
Capital-Utilization Cost Param.	ψ_u	0.656
Quarterly Survival Rate	$\bar{\gamma}$	0.98727
Long-Term Debt Maturation	ρ_b	0.968
S.S. Capital-Income Tax Rate	τ_K	0.218
Interest-Rate Smoothing	ρ_i	0.726
CB Inflation Response (Active)	φ_π^M	2.058
CB Output Response (Active)	φ_y	0.0014

Notes: Preference, labor, growth, production, and depreciation parameters come from Gertler (1999). Price-setting, investment-related, long-term debt, tax, and monetary parameters come from Bianchi et al. (2023). The elasticity ε_p is derived from $\mu_p = \varepsilon_p/(\varepsilon_p - 1)$.

To this end, the model uses $\omega_t = \bar{\omega} - \varepsilon_t^{\text{ret}}$ for the endogenous probability of a worker remaining one in the coming period, where I specify $\varepsilon_t^{\text{ret}} = s_\varepsilon (t/\nu_\varepsilon)^2 \exp(-t/\nu_\varepsilon)$, where t

represents time and s_ε and ν_ε control the magnitude of the shock and its timing, respectively. This gives the retirement wave a hump shape that resembles a Poisson distribution.

After the peak of the wave, ω_t returns to $\bar{\omega}$. I calibrate these so that at its peak, $t = 2\nu_\varepsilon$, with $\nu_\varepsilon = 7.36$ and $s_\varepsilon = 0.00495$. This means the the wave of retirement peaks about 14.7 quarters out into the future. Consequently, the dependency ratio $\psi_{t+1} = (1-\omega_t+\gamma\psi_t)/(1+n_t)$ reaches its zenith of 0.383 approximately 10 years into the future, in accordance with SSA projections. An illustration of number of retirees per worker along with the total social security benefits across the economy is provided in Figure 3. This wave is identical for each of the reforms discussed below.

In the wake of these retirement waves, the government has various options available to it for funding the inexorable expansion in benefits. I focus on two cases: a fully-funded and a fully-unfunded expansion. In the case of a fully funded expansion, the entire benefit expansion is backed by future taxes. Note that this does not mean that each dollar increase in transfers is met by a simultaneous dollar increase in tax revenue. What it *does* mean is that taxes will adjust so as to quickly stabilize the debt-to-GDP ratio. In the case of payroll taxes, for instance, this means $\rho_L = 0.85$ and $\gamma_L = 0.08$ in equation 32.

In the fully unfunded case, we have $\gamma_L^F = 0$, which means that the government makes zero commitment to raise taxes (payroll taxes in this case) to help finance the expansion in benefits and entirely relies upon debt finance. The unfunded expenditure results in an adjustment of the price level through the fiscal-theory mechanism where inflation re-balances the values of outstanding debt and expected tax revenue.

Finally, I simulate four possible paths forward for policy, which I call social security "reforms", for the economy. I compare the four reforms under the fully funded versus fully unfunded case. The first reform is simply no reform at all, meaning the government does nothing to change the program itself. This case represents a raw comparison of the effects of a fully funded versus fully unfunded expansion of the social security program over the next 10 years in order to meet its commitments.

The second reform is an immediate 10% cut to the program for both cohorts (workers and retirees). The cut is not permanent. It is quite persistent, but ultimately recovers back to its baseline. The reason for this specification is that, even if social security is radically transformed so as to generate a 10% cut in benefits overall, the demographic dynamics projected over the next 50—100 years are indicative that, so long as general structure of the current PAYGO program is in place, demographics alone will continue to drive up the overall costs.

The third reform is gradual benefit cut that results in a 1% reduction in the program on impact, and gradually increases over several years before it also, eventually, returns to the pre-cut trend. The fourth and final reform is a future-retiree-only cut, which means that, if the reform is implemented in period t, anyone who is already a retiree in period t retains the same benefits as before the reform. Everyone else who enters retirement at period $t + 1$ an immediate 10% cut relative to the previous benefit level. This cut is gradually reduced as the years advance since the reform is only temporary and should reflect that its persistence depends on the retirement rate.

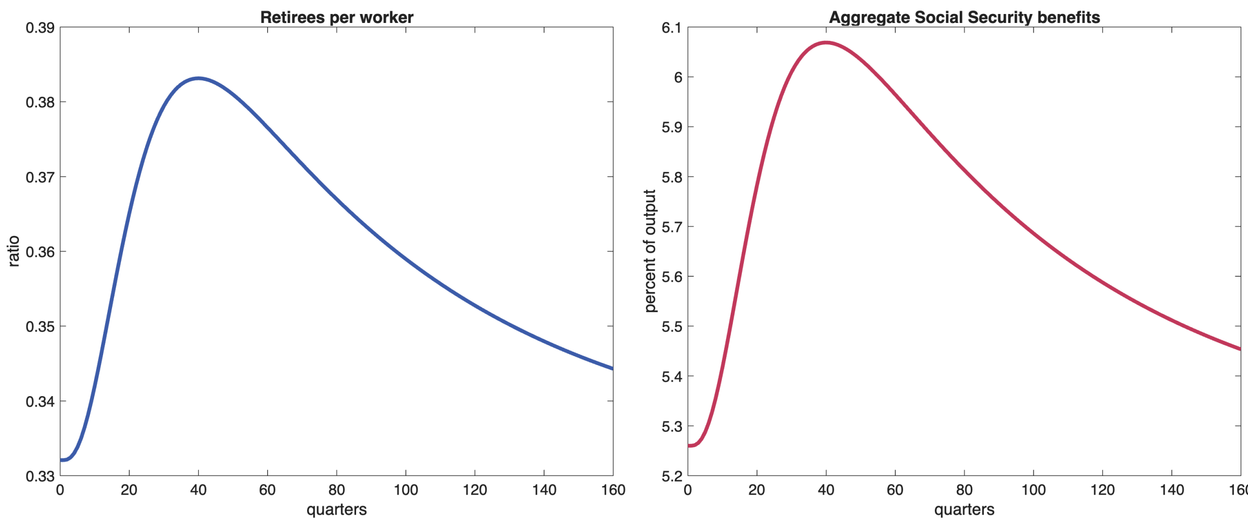

Figure 3: Retirement/Benefit Paths

5 Results

I now turn to the results of the policy experiments. Figure 4 reports the impact at steady state of the ensuing retirement wave captured in Figure 3 among the two different financing regimes (fully funded versus fully unfunded) and four reforms. One of the most striking differences between the two regimes is the magnitude and persistence of inflation on impact following the expansion. Our model predicts that, assuming the government does nothing to reform the social security program and raises no additional future tax revenue to finance the expansion, inflation will rise by 2.85% on impact and remains persistently and substantially above baseline over nearly all horizons. In the fully-funded case, payroll taxes adjust to finance the expenditures, hence the dampened inflation response (with a peak of 0.40% with no reforms) relative to the fully-unfunded case where the inflation adjusts to help stabilize

the additional unfunded debt.

Looking at the debt-to-GDP ratio, it may at first appear that unfunded expansion improves solvency in the no-reform case as debt falls from 101% to 89% of output. However, this effect is purely the result of the inflationary correction of the unfunded transfers: the value of outstanding nominal debt falls, due to its lower real returns, and nominal output rises, pushing the overall ratio down.

One might also wonder why debt-to-GDP rises to 104% under the no-reform "fully-funded" regime. Recall that the payroll tax rule in equation (32) is smoothed and responds to lags of itself and debt. Therefore, taxes to not immediately cover the fiscal expansion and the government must issue long-term debt to cover this financing gap. What "fully-funded" means here is not that debt-to-GDP is held fixed across all horizons, but that the expansion is fully *fiscally backed*, meaning payroll taxes will eventually adjust so as to stabilize and repay the additional debt incurred.

Another noteworthy takeaway is that, in both financing regimes, the gradual and immediate cuts generate macroeconomic effects that are not substantially disparate. Gradual cuts begin with a modest 1% but and then gradually phases in larger cuts before dissapating, while the immediate cut is an all-at-once 10% cut. The immediate cut, by construction, has a larger effect at early horizons, but inflation, investment, and taxes display only negligible differences in their patterns over medium to long-run horizons. Accordingly, it may behoove a social planner to take into consideration the differential welfare effects of an immediate cut, versus gradual cut as the former will lead to drastic changes in consumption patterns among existing retirees in the short run.

The response of investment under the two financing regimes is a little more complicated. Recall that labor is supplied elastically. In the fully-funded case, the benefit cuts give rise

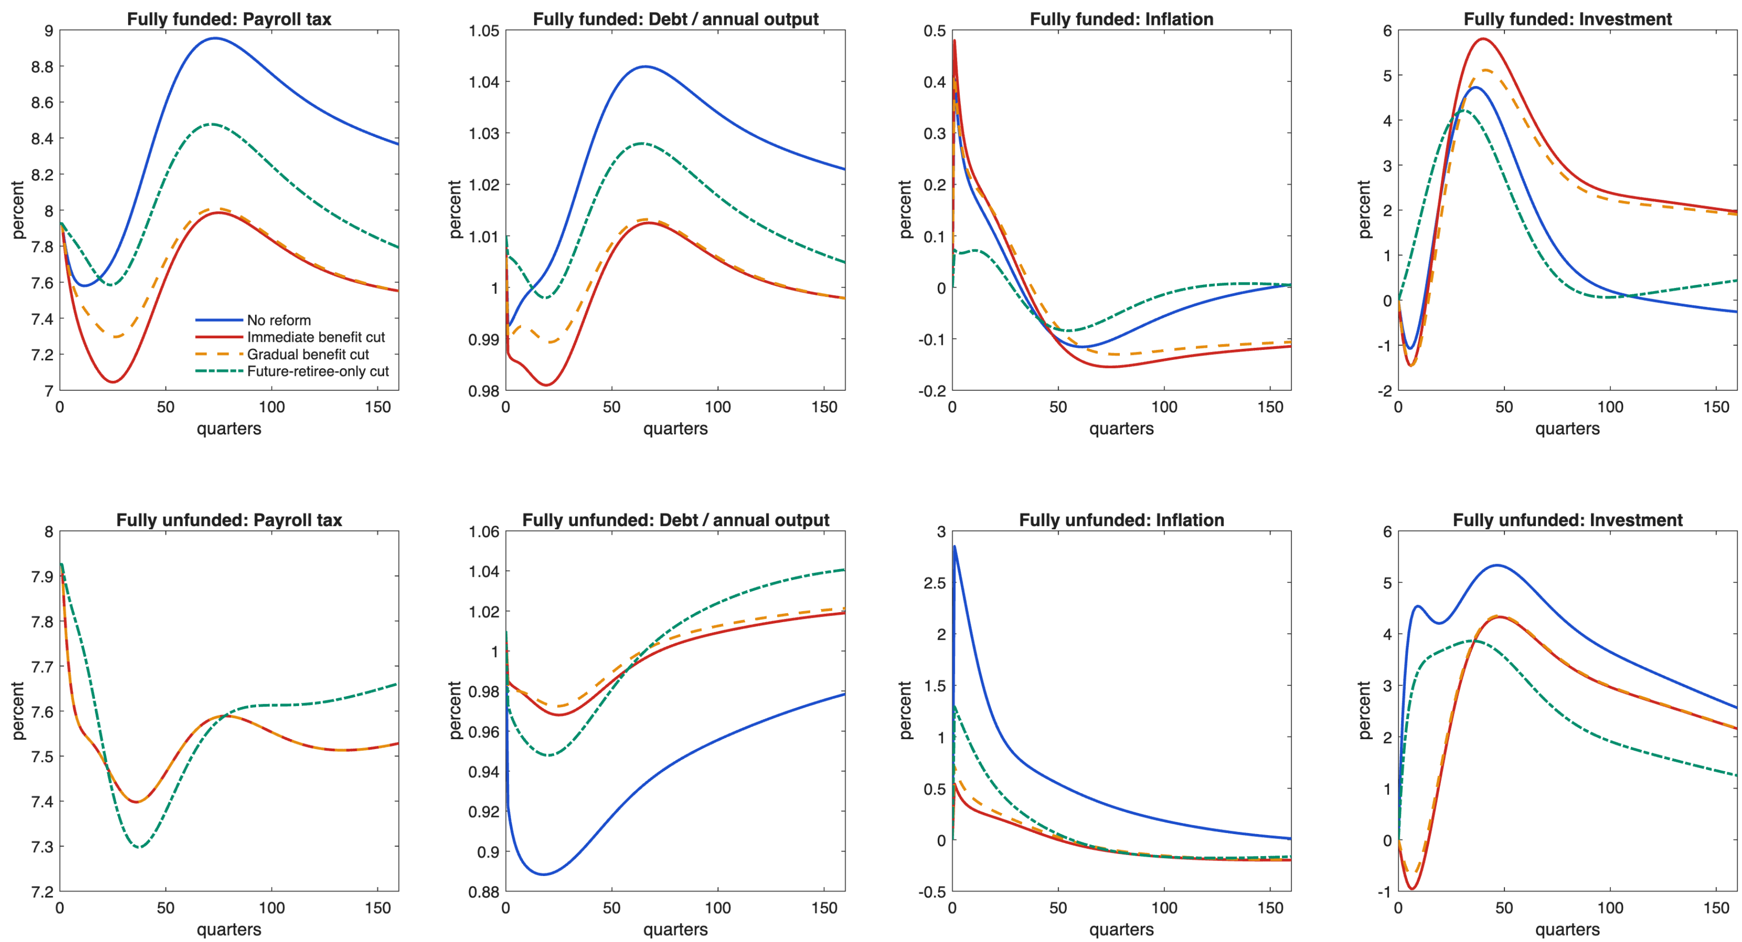

Figure 4: Four Reforms by Financing Regime

The fully-funded regime is in the top row and the full-unfunded regime is in the bottom row. The four reforms are delineated by color. Note that 150 quarters is about 37.5 years. The vertical axis is measured in terms of percentage deviation from steady state.

to lower rates of taxation, which raises the supply of effective labor, eventually yielding a peak investment response of 5.8% in the wake of an immediate benefit cut. Moreover, both workers and retirees operate under perfect foresight. Workers increase consumption today and aggregate demand increases economy-wide, even in the face of the effects on retiree consumption, who have higher MPCs than workers. Thus, under the unfunded no-reform case, for instance, the upward pressure in prices pushes down the real interest rate, while the aggregate demand channel raises the rental rate of capital. These forces combine to raise Tobin's q, leading to an immediate rise in investment. This phenomenon is an artifact of an inflation-driven demand boom. The initial dip in investment across three out of four reforms in the fully-funded case is a byproduct of the monetary policy contraction that follows the rise in inflation.

The next question of interest is, assuming that the government makes the requisite fiscal adjustments to fully fund the expansion to the program, what effects will the effects of taxation be? Specifically, I am interested in how payroll tax financing of the social security expansion affects the economy versus taxes on capital gains. In Figure 5 below, I present results of this analysis. The results represent the fully-funded financing regime. For each IRF pane and reform, I plot the following differential:

$$\Delta_t^{\tau} = (y_{\tau_{K,t}} - \bar{y}_{\tau_K}) - (y_{\tau_{L,t}} - \bar{y}_{\tau_L})$$

Where $\bar{y}_{\tau_j}$ represents the steady state-level of variable y under the steady-state tax regime on the object of interest $j \in \{K, L\}$. For analytical purposes, I treat each tax instrument as mutually exclusive: agents face either payroll taxes or capital taxes, not both. Positive movements in the IRFs below correspond to cases where the outcome variable under capital taxation exceeds its level at that particular horizon if, instead, the social security expansion had been financed via payroll taxes.

The results indicate that, absent any reform of the pension program, capital taxation leads to an increase in effective labor, which eventually peaks at 0.51%, but lowers investment by about 2.46%, the capital stock by 1.61%, and overall GDP by about 0.58%. On net, then, capital taxation's effects on investment and output overwhelm the gains wrought by the increase in effective labor. This pattern can initially be reversed by cuts to the program, however. The reason for this is owing to the fact that immediate spending cuts beget tax cuts. Moreover, capital tax cuts raise the returns on capital, Tobin's q, and, therefore, investment overall. Relative to a payroll tax cut, the capital tax cuts generate an increase in investment by as much as 1.89% and output by 0.57% following the immediate 10% reduction in social security outlays.

Future-retiree-only reform yields a kind of middle ground. Over immediate and medium-run horizons, the IRFs behave in a manner akin to the no-reform baseline in that capital taxation contracts activity relative to payroll taxation. This is due to the fact that existing retirees continue to enjoy their previous benefit levels as long as they remain alive, requiring continued financing of these old benefits. It is only through individuals entering retirement *after* the reform that fiscal adjustments begin to accumulate, which eventually leads this reform to converge in its effects with the other reforms to the program after three and a half decades.

In terms of actual debt accumulation, the differences between the reforms are relatively small. The largest differential is about 1.8% of annual output. Since all of these cases represent full fiscal backing of the expansion, the choice of tax instrument largely affects the timing of adjustments to the government's debt dynamics. Overall, the choice of tax instrument is consequential. Absent any reform to the social security program, financing its expansion via capital income taxes reduces investment, the capital stock and, moreover, output. These effects can be reversed if cuts are made to the program due to the ability to simultaneously reduce the tax burden. The strongest stimulus come from capital taxation relief in this case. Output and investment respond more strongly the more rapidly the cuts are implemented. Given the fact that differences in the overall magnitude of debt-to-GDP are fairly negligible across reforms (including no reform), the primary consideration in implementing a given reform ought to be the effects it has on private activity, given the model's predictions.

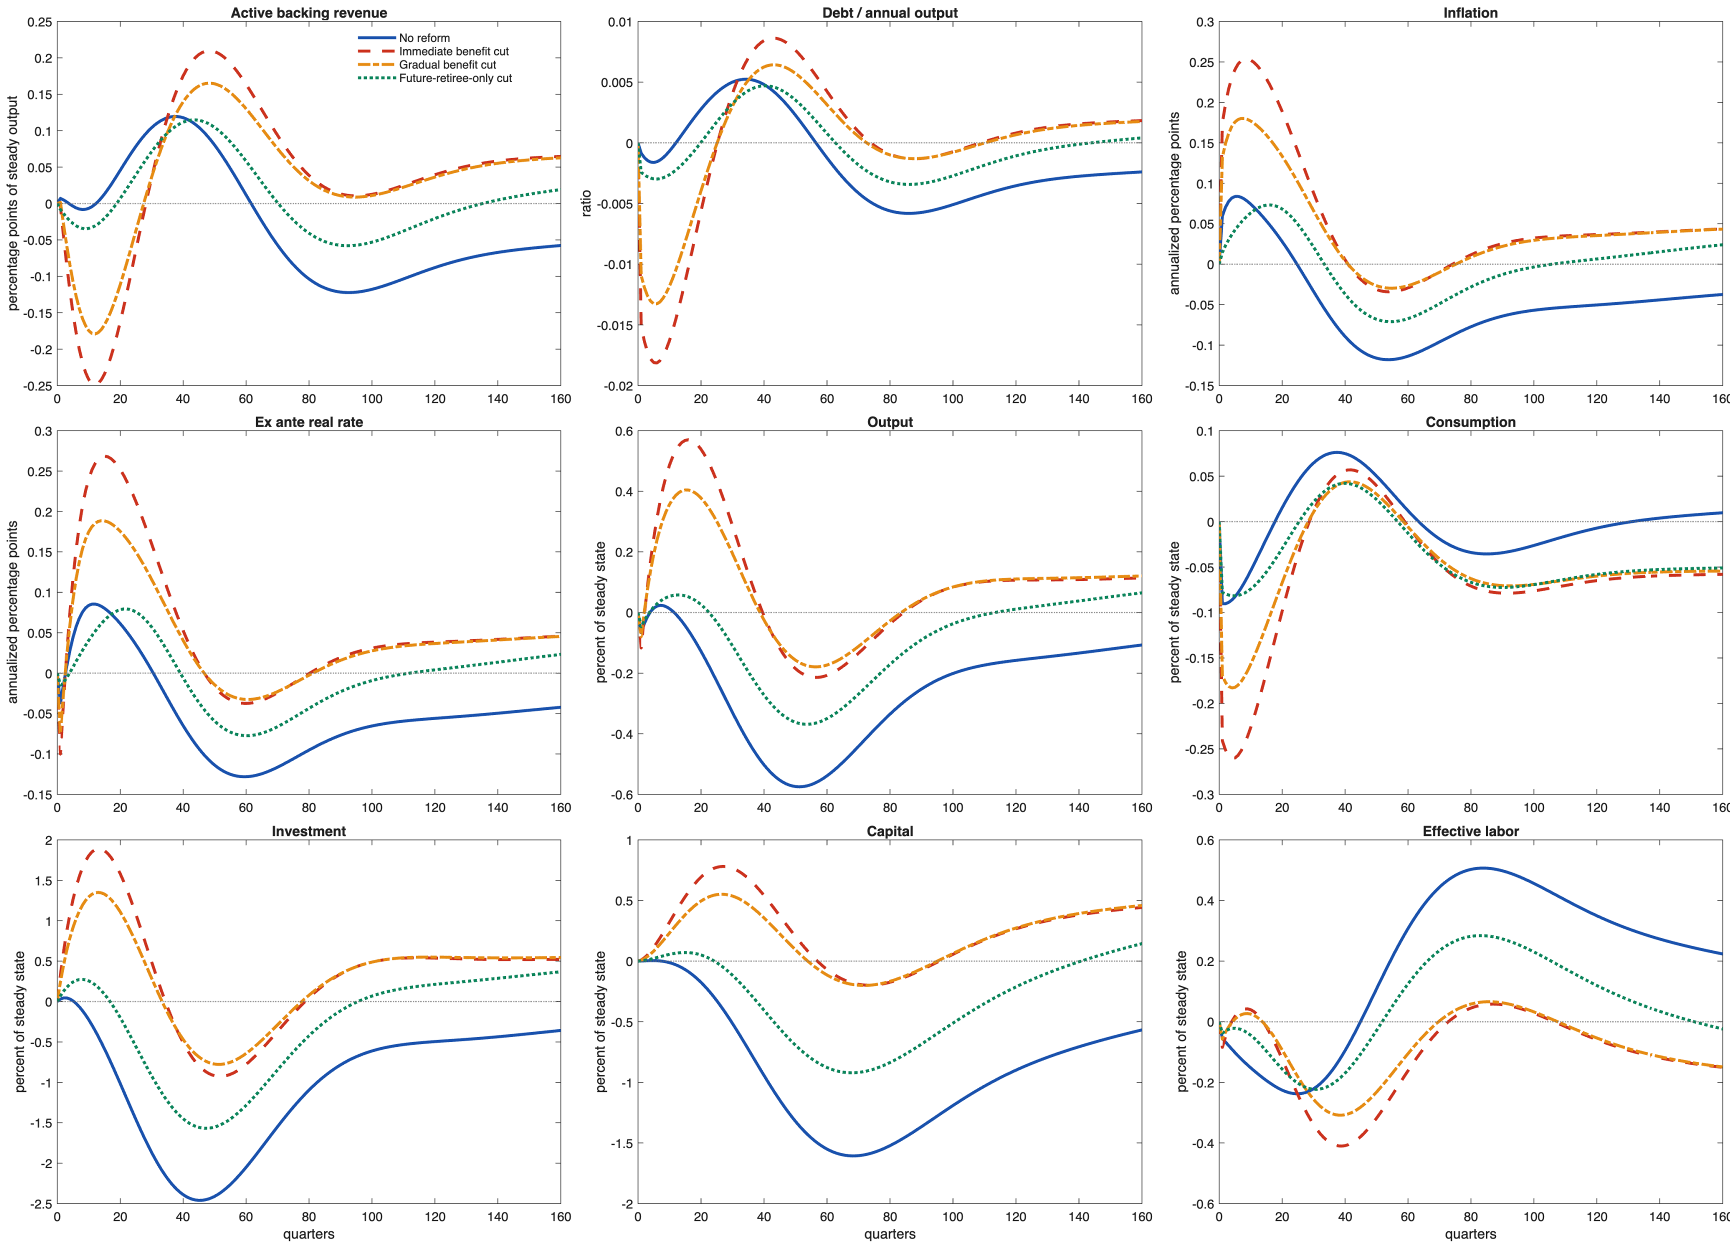

Figure 5: Four Reforms by Tax Scheme

Each IRF, horizon by horizon, is the differential response of a given variable under a capital-income tax relative a payroll tax. The four reforms are delineated by color. Note that 150 quarters is about 37.5 years. The vertical axis is measured in terms of percentage deviation from steady state.

6 Welfare Analysis

At this point, I perform a welfare analysis to determine, among the menu of policy options (unfunded vs funded; cuts vs no reform), which combination of choices is welfare maximizing. To this end, I use the same household preferences from Gertler (1999) to evaluate how debt, inflation, taxes and asset price variation affects household consumption and leisure. The form of preferences employed by Gertler is taken from Farmer (1990), which the latter calls "risk neutrality and constant elasticity of substitution" or "RINCE" preferences. The point of these preferences is to make agents risk-neutral with respect to income risk. This is important in recreating life-cycle behavior because it prevents workers from engaging in unrealistic additional precautionary saving for retirement "risk." Workers do make preparations for retirement through saving, given the fact that their labor income vanishes and overall income falls as a result, but the exogenous probability of retirement baked into the model to generate life-cycle dynamics doesn't additionally create an element of risk that itself induces precautionary saving.

I follow a standard consumption-equivalent approach to the welfare analysis in the spirit of Lucas (1990), where the relevant metric of interest in measuring welfare is, ultimately differences in lifetime consumption; specifically, in our context, the differences in consumption that would make households indifferent between different fiscal regimes under the various reforms. The CEV approach is modified to account for the life-cycle features of the model in a similar manner to Peterman and Sager (2022).

6.1 Consumption Equivalent Variation

In the model, agents derive utility from consumption and leisure according to $c_t^{\nu}(1-\ell_t)^{1-\nu}$. Denote the consumption-equivalent variation (CEV) by λ so that we have:

$$c_t^{\nu}(1-\ell_t)^{1-\nu} = [(1+\lambda)c_t^B]^{\nu}(1-\ell_t^B)^{1-\nu}$$

Where c_t^B represents consumption under our benchmark policy that we compare all other policy prescriptions to. Simplifying the above expression we see that:

$$c_t^{\nu}(1-\ell_t)^{1-\nu} = (1+\lambda)^{\nu}\left(c_t^B\right)^{\nu}\left(1-\ell_t^B\right)^{1-\nu}$$

Which shows us that λ represents the requisite percentage change of consumption in the benchmark case necessary to make an agent indifferent between the benchmark policy and the alternative policy. Now, following Gertler (1999), I guess the following solutions to the value recursions for workers and retirees in section 2.1:

$$\begin{aligned}
V_t^{w,P} &= (\mu_t^P)^{-1/\rho} c_t^{w,P}\left[\frac{\varsigma}{(1-\tau_t^{i,P})w_t^P}\right]^{1-\nu} \\
V_t^{r,P} &= (\epsilon_t \mu_t^P)^{-1/\rho} c_t^{r,P}\left[\frac{\varsigma}{(1-\tau_t^{i,P})\xi w_t^P}\right]^{1-\nu}
\end{aligned}$$

where $V_t^{w,P}$ represents the value function of the worker w under policy P. Since the value functions (7) and (8) have the same homogeneity properties as the consumption-leisure composite above, we can write:

$$V_{g,s}\left((1+\lambda)c^B, \ell^B\right) = (1+\lambda)^{v} V_{g,s}^B$$

Where $V_{g,s}^B$ is the lifetime utility of group g (workers, retirees), in state s under the benchmark policy B. By the definition of CEV, we have:

$$\begin{aligned}
V_{g,s}^P &= V_{g,s}\left((1+\lambda)c^B, \ell^B\right) \\
&= (1+\lambda)^{v} V_{g,s}^B \\
&\Rightarrow \lambda = \left(\frac{V_{g,s}^P}{V_{g,s}^B}\right)^{1/\nu} - 1
\end{aligned}$$

The CEV that is actually reported takes into account contemporary and future cohorts of individuals. I demarcate the overall welfare of those alive at announcement date of the reform (date 0) according to:

$$W_0 = V_0^{w,P} + \psi_0 V_0^{r,P}$$

With $V_0\left(\{(1+\lambda)c_\tau^B, \ell_\tau^B\}_{\tau=t_0}^{\infty}\right) = V_0^P$. The welfare analysis also takes into account entrants e in the form of workers with no financial assets, so that their consumption at date t is characterized by $c_t^{e,P} = \mu_t^{w,P}(H_t^{w,P} + S_t^{w,P})$, with the lifetime value of $V^{e,P}$ under policy P. The purpose of including entrants is to prevent the population from being static and to analyze the effects of the various policy proposals on future generations who will be contributing tax dollars to finance the pension program. Therefore, the ultimate welfare object of interest under policy P is:

$$\mathbb{W}^P = W_0 + \sum_{t=1}^{\infty} \vartheta^t \gamma_x^{\nu t} M_t V_t^{e,P}$$

M_t represents new entrants into the economy at time t that replace (and expand upon) the new retirees. It is given by $M_{t+1} = N_{t+1}^w - \omega N_t^w = (1+n-\omega)N_t^w$, all of which follows from equations 1 and 2. The parameter γ_x accounts for technological growth's effects on future consumption, while ϑ discounts future generations relative to living cohorts in the calculation of overall welfare. This is necessary because, one, we genuinely do place a higher weight on the effects a policy has on living generations, relative to future ones, when considering the utility of that policy; and two, we need the second term in $\mathbb{W}^P$ to converge, which means we need: $\vartheta\gamma_x^{\nu}(1+n) < 1$. In terms of the actual calibration of ϑ, its quarterly value is set according to $\vartheta = (1+r_s)^{-t/4}$, where I calibrate $r_s = 0.02$, so that 40 years in the future, entrants receive 0.453 times the weight as workers entering today.

This ultimately enables us to construct the CEV of policy P, which is what is reported in the accompanying tables of this study:

$$CEV^P = 100\left[\left(\frac{W^P}{W^B}\right)^{1/\nu} - 1\right]$$

6.2 Welfare Results

Below, in Table 3, are the results of the welfare analysis vis-à-vis the four reforms and their respective financing regimes. Every cell compares a particular reform against the baseline case of no reform under the same financing regime. The table indicates what the incremental effect is of a policy reform while holding the financing regime fixed.

The table tells an intuitive story: fully-funded immediate and gradual social security cuts result in shifts of welfare from the average incumbent retiree towards workers. Future-retiree-only cuts keep the average incumbent retiree whole, while reducing welfare of the average worker and early entrants. Under the fully-unfunded BFM regime, living cohorts gain while future entrants bear the subsequent burden. Note that whatever tax is adjusted is the sole tax adjusted in that case; the other is left unchanged. Steady state tax rates are calibrated so that the baseline payroll tax rate is 7.93%,[5] and the capital income tax rate is 21.8%, the latter of which is taken from Leeper et al. (2017). In every case, the benefit path is identical. When each respective tax is adjusted, it is adjusted to create the same revenue effect. Therefore, the welfare comparison is intended to convey the effects of changes in the tax base within the same financing regime, whereas differences in financing regimes unveil effects of fiscal backing strength.

Table 3: Welfare Effects of Benefit Reform By Financing Regime

Reform	Financeing Regime	Worker	Retiree	Living	Entrant 0	Future Entr.	All
		Announcement-date lifetime CEV				CEV at $r_s = 2\%$	
Immediate benefit cut	Funded: payroll-tax adjustment	1.108	-1.775	0.165	1.216	1.017	0.507
Gradual benefit cut	Funded: payroll-tax adjustment	0.845	-1.136	0.198	0.695	0.990	0.516
Future-retiree-only cut	Funded: payroll-tax adjustment	-0.328	1.636	0.309	-0.784	0.734	0.480
Immediate benefit cut	Funded: capital-income-tax adjustment	1.905	-0.966	0.966	1.068	0.258	0.680
Gradual benefit cut	Funded: capital-income-tax adjustment	1.433	-0.531	0.791	0.493	0.251	0.573
Future-retiree-only cut	Funded: capital-income-tax adjustment	-0.054	1.918	0.585	-0.927	0.091	0.386
Immediate benefit cut	Unfunded: no tax adjustment	1.475	0.655	1.209	-2.870	-0.892	0.355
Gradual benefit cut	Unfunded: no tax adjustment	1.198	1.024	1.142	-2.774	-0.920	0.304
Future-retiree-only cut	Unfunded: no tax adjustment	-0.132	2.872	0.834	-2.694	-1.020	0.081

Notes: Entries are percent permanent consumption equivalents relative to no reform under the same financing rule and at the same date. Positive values favor the reform. No reform includes the common retirement wave. Worker and retiree entries are lifecycle-status averages, not birth cohorts. "Living" denotes workers + retirees living at time of reform implementation. Entrant 0 is a zero-financial-asset worker entering at announcement. Future-entrant ("Future Entr.") and all-generation ("All") columns use explicit equal-person weights and an annual social discount rate $r_s = 2\%$.

[5] This is model-implied by dividing the OASDI benefit cost as a percent of GDP by the steady-state payroll base.

Under two of the three financing regimes: fully-unfunded and funded with capital-income-tax adjustment, the immediate benefit cut generates the highest CEV for all generations. In the third case, fully funded with payroll-tax adjustment, the gradual benefit cut generates the highest CEV. In the case of the immediate cut under a fully-funded-capital-tax adjustment, the average worker alive at the time of the policy announcement values the policy as much as a 1.905% increase in consumption for the remainder of their life under the no-reform baseline of that same financing regime. For the retiree, by contrast, they would be willing to forfeit up to 0.966% of their consumption, permanently, to avoid suffering the policy change and to remain in the no-reform state. Note that under the "Living" column is the CEV of those alive when the reform is announced, which includes both workers and retirees and is not a simple sum of worker and retiree values. The combined CEVs under the Living category are weighted by the respective population shares of workers and retirees, where workers make up approximately 67.5% of the population.

In the appendix is displayed Table 5, which compares each reform to the same baseline: the fully-funded payroll-adjusted no-reform case and ranks each reform under each financing regime by its overall CEV of everyone alive at the time of the announcement. This construction favors an immediate cut under a fully-funded regime with a corresponding capital-income tax. The difference in welfare between this case and the next best policy is significant; a roughly 31% improvement. Additionally, under this construction, the worst-case scenario, in terms of lifetime welfare of those alive at the time of the announcement, is the no-reform fully-unfunded case, where those alive would be indifferent to a 1.108% cut to lifetime consumption under the fully-funded payroll-tax baseline and the unfunded no-reform case.

Table 4 is intended to complement the tax-instrument IRFs in Figure 5. Each entry directly compares the two funded marginal tax instruments for the same benefit rule.

In the case of benefit cuts, the comparison is between tax *cuts* that accompany the reform of social security benefits. The capital-income-tax adjustment raises the welfare of all generations (present and discounted future generations) relative to payroll adjustment for all four benefit rules. The no-reform case demonstrates that the results that favor capital tax adjustments only favor that form of tax adjustment due because they are cuts. If, on the other hand, taxes need to be *hiked* to finance the expansion in benefits, taxing capital reduces welfare across all living individuals at the time of the reform. This happens because living workers and retirees consume according to:

$$C_t^w = \mu_t^w \left[(1-\lambda_t)F_t + H_t^w + S_t^w\right]$$
$$C_t^r = \epsilon_t \mu_t^w \left[\lambda_t F_t + H_t^r + S_t^r\right]$$

Where F_t is income derived from ownership of a financial portfolio that includes claims to capital and bonds. In the no-reform case, the tax hike is announced before it is actually implemented, since the fiscal rule exhibits inertia. In the case of a capital-income tax, this affects asset valuations immediately as asset prices are forward-looking. Moreover, over the transition after the initial shock, capital continues to fall up to 2.8% below its steady state, which leads living agents to prefer payroll-tax financing of the benefit expansion.

This also explains why future entrants prefer capital-tax to payroll-tax financing of the expansion in the no-reform case: entrants start out with zero assets. Therefore, they are exempt from the immediate revaluation effect on the asset portfolio of the announced tax hike. Furthermore, relative to the payroll-tax hike, future entrants enjoy higher lifetime consumption and leisure under the capital-financed hike. Therefore, the majority of the incidence on capital-tax-financed expansions to the program falls on those currently living.

Table 4: Capital-Income-Tax Relative to Payroll-Tax Adjustments

Benefit rule	Worker	Retiree	All alive	Entrant 0	Future entrants	All generations
	Announcement-date lifetime CEV				CEV at $r_s = 2\%$	
No reform	-0.232	-0.215	-0.226	0.336	0.677	0.136
Immediate benefit cut	0.555	0.606	0.572	0.189	-0.080	0.309
Gradual benefit cut	0.351	0.396	0.365	0.135	-0.060	0.194
Future-retiree-only cut	0.042	0.062	0.048	0.192	0.034	0.043

Notes: Each cell directly compares capital-income-tax adjustment with payroll-tax adjustment for the same benefit rule. Positive values favor capital-income-tax adjustment. Both experiments retain both calibrated baseline tax rates; the labels identify the marginal backing instrument.

7 Conclusion

This paper incorporates new modeling tools into an established framework for analyzing social security policy in the United States. In particular, I am able to assess the macroeconomic and welfare consequences of four types of reforms to the social security program across two kinds of financing regimes: fully-funded versus fully unfunded. This latter distinction is especially pertinent, given the current state and trajectories of U.S. government debt. Using a standard measure of welfare, I find that the worst policy path forward for currently-living generations is our current trajectory: unfunded fiscal expansion to finance social security entitlements with no reform to the program. This path has the effect of generating large fiscal inflation, creating price distortions, and significant portfolio effects through the revaluation of government debt. It also leads to less leisure for both workers and existing retirees. When taking into account both present and future generations' welfare, neglecting to reform the program while fully funding its expansion via payroll taxes is least recommended. In all cases, an immediate 10% cut to the program that generates the requisite fiscal space to lower tax rate on capital income performs best. Taken as a whole, this approach generally recommends a commitment to full-funding of future obligations over and above unfunded expansions.

References

Barro, R. J. (1974, November). Are Government Bonds Net Wealth? *Journal of Political Economy 82*(6), 1095–1117.

Bianchi, F., R. Faccini, and L. Melosi (2023, November). A fiscal theory of persistent inflation. *The Quarterly Journal of Economics 138*(4), 2127–2179.

Blanchard, O. J. (1985). Debt, Deficits, and Finite Horizons. *Journal of Political Economy 93*(2), 223–247.

Carvalho, C., A. Ferrero, and F. Nechio (2016). Demographics and real interest rates: Inspecting the mechanism. *European Economic Review 88*, 208–226.

Cochrane, J. H. (2001, January). Long-Term Debt and Optimal Policy in the Fiscal Theory of the Price Level. *Econometrica 69*(1), 69–116.

Cochrane, J. H. (2022). The fiscal roots of inflation. *Review of Economic Dynamics 45*, 22–40.

Cochrane, J. H. (2023). *The Fiscal Theory of the Price Level*. Princeton: Princeton University Press.

Farmer, R. E. A. (1990, February). RINCE Preferences. *The Quarterly Journal of Economics 105*(1), 43–60.

Fujiwara, I. and Y. Teranishi (2008). A dynamic new Keynesian life-cycle model: Societal aging, demographics, and monetary policy. *Journal of Economic Dynamics and Control 32*(8), 2398–2427.

Gertler, M. (1999). Government debt and social security in a life-cycle economy. *Carnegie-Rochester Conference Series on Public Policy 50*, 61–110.

Hosseini, R. and A. Shourideh (2019). Retirement Financing: An Optimal Reform Approach. *Econometrica 87*(4), 1205–1265.

Kara, E. and L. Von Thadden (2010). Interest Rate Effects of Demographic Changes in a New-Keynesian Life-Cycle Framework. *SSRN Electronic Journal*.

Leeper, E. M. (1991, February). Equilibria under 'active' and 'passive' monetary and fiscal policies. *Journal of Monetary Economics 27*(1), 129–147.

Leeper, E. M., N. Traum, and T. B. Walker (2017, August). Clearing Up the Fiscal Multiplier Morass. *American Economic Review 107*(8), 2409–2454.

Lucas, R. E. (1990). *Models of Business Cycles* (Repr. d. Paperback-Ausg. von 1988 ed.). Yrjö Jahnsson Lectures. Oxford: Blackwell.

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

Peterman, W. B. and E. Sager (2022, October). Optimal Public Debt with Life Cycle Motives. *American Economic Journal: Macroeconomics 14*(4), 404–437.

Romer, D. (1988, January). What Are the Costs of Excessive Deficits? *NBER Macroeconomics Annual 3*, 63–98.

Sargent, T. J. (1981, September). Some Unpleasant Monetarist Arithmetic. *Quarterly Review 5*(3).

Schmitt-Grohé, S. and M. Uribe (2000, February). Price level determinacy and monetary policy under a balanced-budget requirement. *Journal of Monetary Economics 45*(1), 211–246.

Sims, C. A. (1994, May). A simple model for study of the determination of the price level and the interaction of monetary and fiscal policy. *Economic Theory 4*(3), 381–399.

Sims, C. A. (2011). Stepping on a rake: The role of fiscal policy in the inflation of the 1970s. *European Economic Review 55*(1), 48–56.

Weil, P. (1989). Overlapping families of infinitely-lived agents. *Journal of Public Economics 38*(2), 183–198.

Woodford, M. (2011, December). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.

A Model Details

A.1 Capital Accumulation and Firm Value

A.1.1 Investment adjustment costs

Let investment growth be denoted by:

$$x_t \equiv \frac{I_t}{I_{t-1}}$$

Let stationary investment be I_t and define the installation-loss function as:

$$S(x_t) = \frac{\phi_I}{2} (x_t - 1)^2 ,$$

Installed capital evolves according to:

$$k_{t+1} = (1-\delta)k_t + [1 - S(x_t)]I_t$$

A capital manager chooses investment and capital utilization to maximize its market value. Let q_t be the shadow value of installed capital. The investment first-order condition is:

$$\begin{aligned} 1 = q_t \left[1 - S(x_t) - S'(x_t)x_t\right] \\ + \frac{q_{t+1}}{R_{t+1}} S'(x_{t+1})x_{t+1}^2 \end{aligned}$$

With the installed-capital equilibrium return condition:

$$R_{t+1}q_t = (1-\tau_{t+1}^K)r_{t+1}^k u_{t+1} - a(u_{t+1}) + (1-\delta)q_{t+1}$$

Capital-manager dividends are

$$D_t^K = [(1-\tau_t^K)r_t^k u_t - a(u_t)]k_t - I_t$$

These conditions follow Bianchi et al. (2023) Section III.A.1 and Appendix equations (D.9)–(D.13). Equations are non-linear counterparts of their capital block. The ownership of the capital manager by Gertler's mutual fund is assumed.

B Welfare Analysis

B.1 Ranking of Packages

Table 5 uses the benchmark case of a regime that is fully funded and payroll tax financed under a no reform benefit rule at the specified two-percent discount rate on future generations criterion to compare various policy package proposals and rank them by corresponding CEV relative to the benchmark case. The ranking complements data in Table 3 and 6.

Table 5: Policy-Package Ranking Benchmark

Rank	Benefit rule	Marginal financing	All alive CEV	Future entrants	All generations
				CEV at $r_s = 2\%$	
1	Immediate benefit cut	Funded: capital-income-tax adjustment	0.737	0.936	0.817
2	Gradual benefit cut	Funded: capital-income-tax adjustment	0.563	0.929	0.710
3	Future-retiree-only cut	Funded: capital-income-tax adjustment	0.357	0.768	0.523
4	Gradual benefit cut	Funded: payroll-tax adjustment	0.198	0.990	0.516
5	Immediate benefit cut	Funded: payroll-tax adjustment	0.165	1.017	0.507
6	Future-retiree-only cut	Funded: payroll-tax adjustment	0.309	0.734	0.480
7	Immediate benefit cut	Unfunded: no tax adjustment	0.088	0.941	0.430
8	Gradual benefit cut	Unfunded: no tax adjustment	0.021	0.913	0.379
9	Future-retiree-only cut	Unfunded: no tax adjustment	-0.283	0.811	0.156
10	No reform	Funded: capital-income-tax adjustment	-0.226	0.677	0.136
11	No reform	Unfunded: no tax adjustment	-1.108	1.850	0.075
12	No reform	Funded: payroll-tax adjustment	0.000	0.000	0.000

Notes: The common benchmark is fully funded no reform with payroll-tax adjustment, evaluated at the announcement. Rankings use unrounded all-generation welfare. The all-generation ranking is conditional on equal-person cardinal utility aggregation and $r_s = 2\%$; it is not a model-implied Pareto ordering.

The immediate capital-tax-backed cut is ranked first and the gradual capital-tax-backed cut second. Their positions are robust to changes in the social-discount factor, which can be seen in Table 7.

B.2 Effects on Future Worker-Entrant

Table 6 allows provides information about what happens to a given worker's lifetime welfare, depending on, one, which reform package was chosen and, two, when they entered the economy, relative to the announcement date 0. The paths are measured against no reform for an entrant at the same date and under the same financing regime.

Table 6: Welfare Incidence Across Future Worker-Entrant Cohorts

Benefit reform	Marginal financing	Year 0	Year 5	Year 10	Year 20	Year 40	Minimum 0–40 years	Year of min.	Mass losing (%)
		Entrant lifetime CEV							
Immediate benefit cut	Funded: payroll-tax adjustment	1.216	1.911	1.883	1.547	1.138	1.138	40	0.0
Gradual benefit cut	Funded: payroll-tax adjustment	0.695	1.682	1.919	1.573	1.160	0.695	0.0	0.0
Future-retiree-only cut	Funded: payroll-tax adjustment	-0.784	-0.019	0.647	1.156	1.376	-0.784	0.0	15.7
Immediate benefit cut	Funded: capital-income-tax adjustment	1.068	0.678	0.461	0.383	0.191	0.191	40	0.0
Gradual benefit cut	Funded: capital-income-tax adjustment	0.493	0.633	0.609	0.411	0.206	0.206	40	0.0
Future-retiree-only cut	Funded: capital-income-tax adjustment	-0.927	-0.527	-0.084	0.245	0.405	-0.927	0.0	34.0
Immediate benefit cut	Unfunded: no tax adjustment	-2.870	-1.772	-1.649	-1.336	-0.886	-2.870	0.0	100
Gradual benefit cut	Unfunded: no tax adjustment	-2.774	-1.819	-1.700	-1.388	-0.925	-2.774	0.0	100
Future-retiree-only cut	Unfunded: no tax adjustment	-2.694	-2.001	-1.832	-1.543	-1.053	-2.694	0.0	100

Notes: Each entry compares a zero-financial-asset worker entering at the stated date with an entrant at the same date under no reform and the same financing regime. "Minimum" is the lowest CEV observed of an entrant in a quarter from years 0 through 40. "Year of min." tells us when the cohort that is most adversely affected enters the economy. Finally, "Mass losing" gives us the population-weighted percent of new entrants between quarters 1 and 160 whose CEV is negative.

Benefit cuts accompanied by tax adjustment generally benefit future worker entrants. On the other hand, grandfathering current retirees shifts losses onto cohorts entering near the reform date. Importantly, under the zero-funded regime, all new-worker cohorts over the first 40 years lose, which indicates a significant redistribution away from future entrants towards generations alive when the reform is announced.

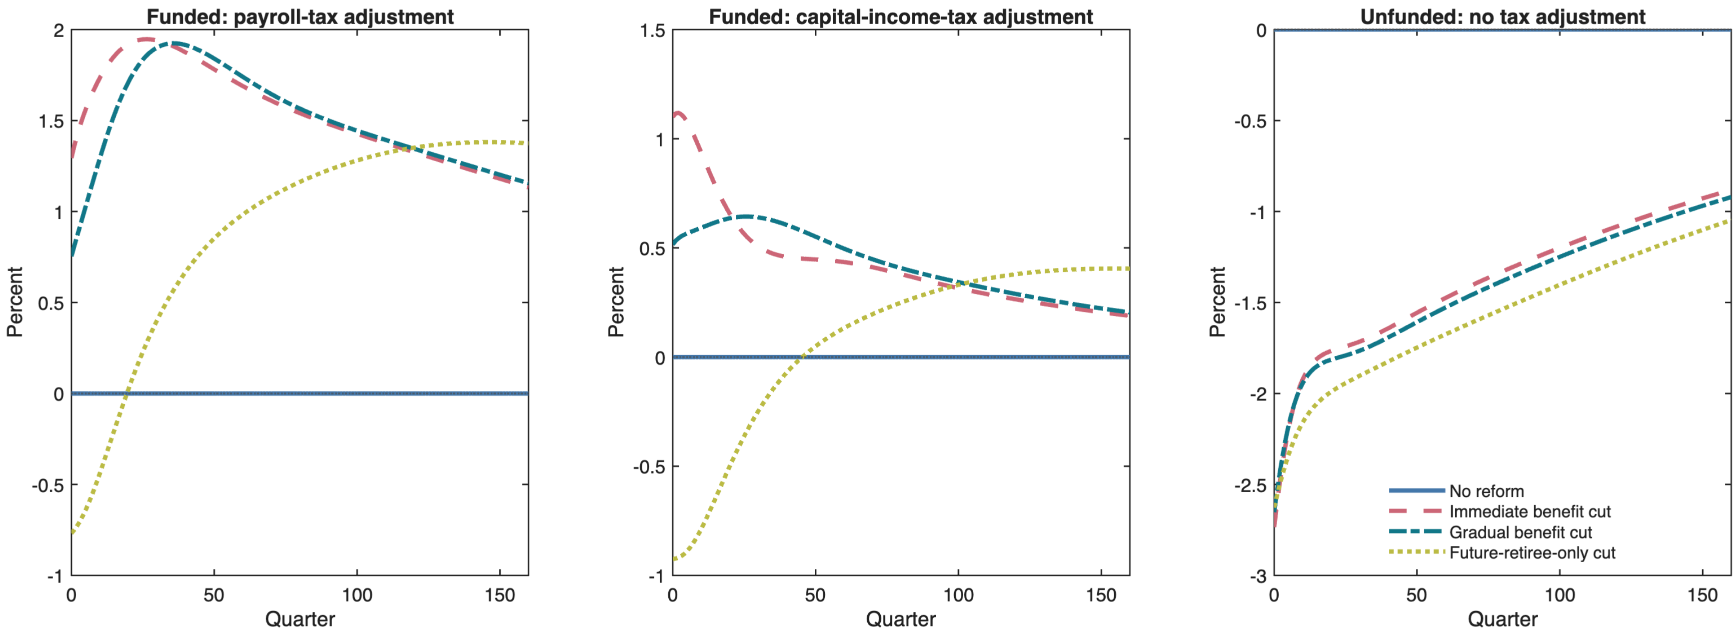

Figure 6: Welfare by worker-entry date.

Notes: Each line represents the *lifetime* consumption-equivalent welfare effect, in percent, for the worker who enters in the indicated quarter on the horizontal axis following the announcement of the policy change. Note that the entrant enters with zero financial assets. Each reform is compared with the no-reform baseline for a worker entering in the same quarter under the same financing rule. Positive values favor the reform; negative values favor the no-reform baseline. Thus, a value of 1 means that consumption under no reform would have to be permanently 1 percent higher to make the entrant indifferent between the two allocations. The three panels hold the financing scheme fixed at payroll-tax adjustment, capital-income-tax adjustment, or no tax adjustment, respectively.

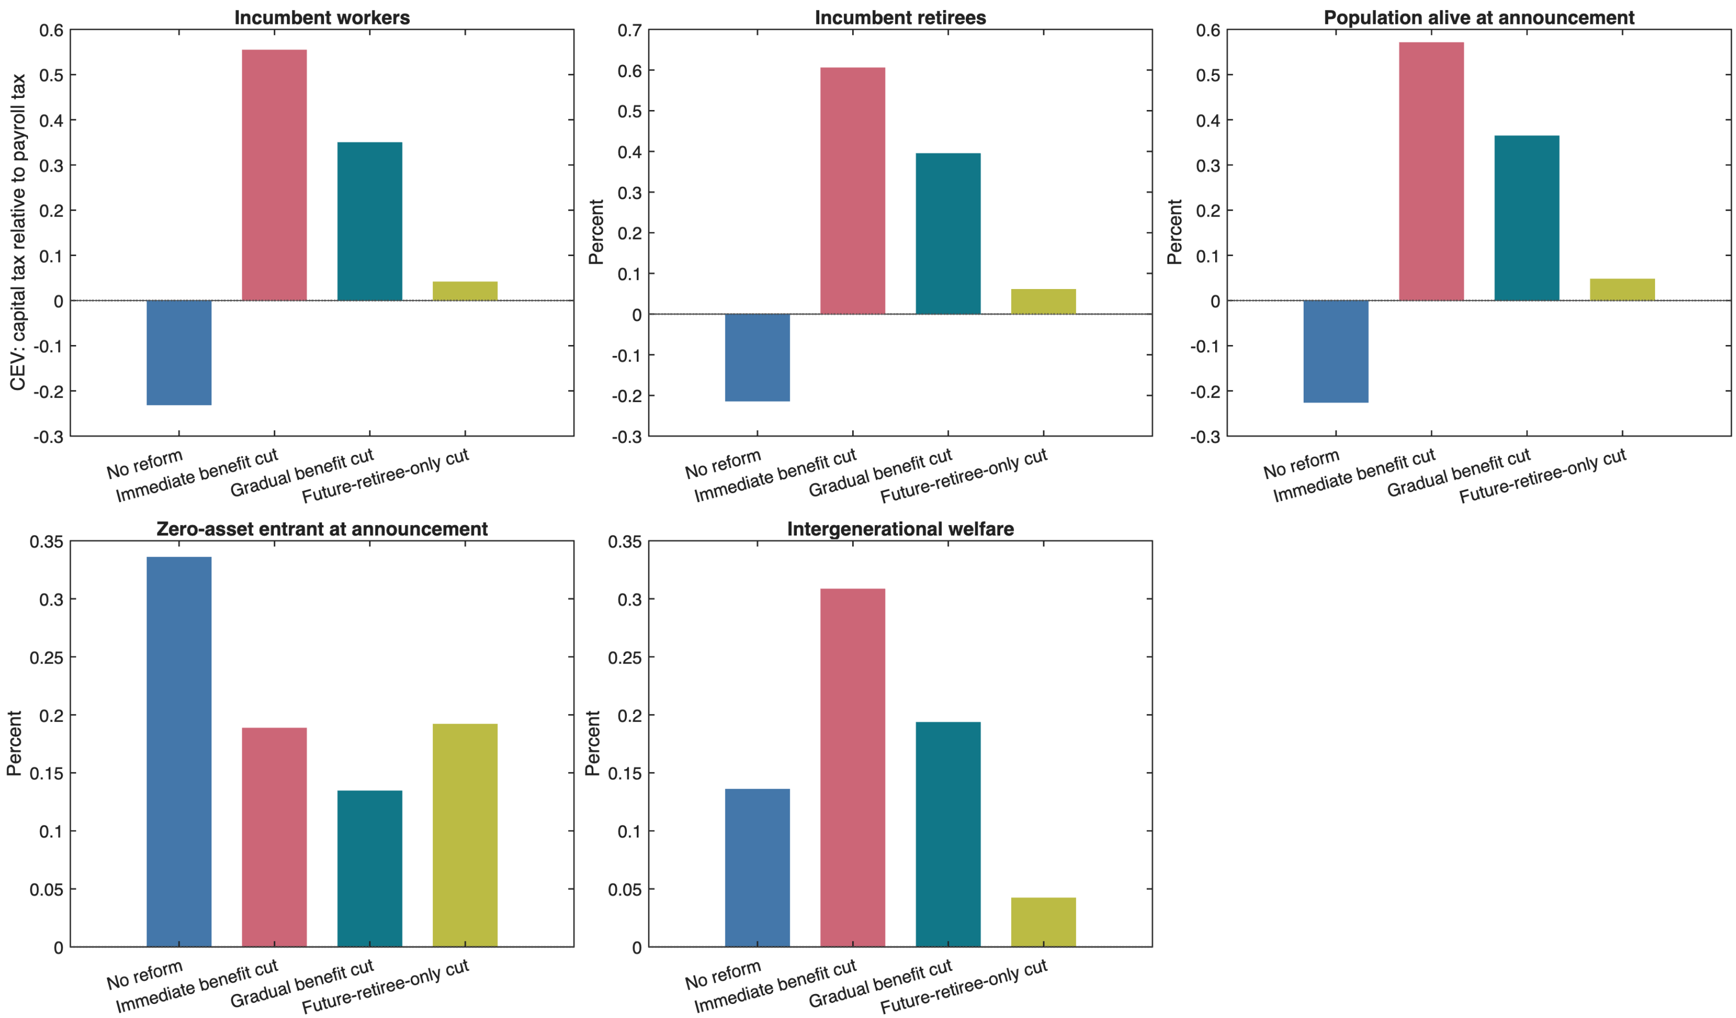

Figure 7: Welfare effects of capital-income-tax adjustment relative to payroll-tax adjustment.

Notes: Each bar reports the consumption-equivalent welfare change, in percent, from using capital-income-tax adjustment instead of payroll-tax adjustment for the same benefit rule. Positive values favor capital-tax adjustment, whereas negative values favor payroll-tax adjustment. Both economies have the same baseline tax rates. Intergenerational welfare uses the two percent annual social discount rate discussed in section 6.1.

B.3 Sensitivity to Social Weights

Table 7 keeps current people, future entrants, and alternative social discount rates visible. The final two columns provide the "Best" and "Worst" rank of each reform over a range of values of social discount rates (0.75–5%).

Table 7: Sensitivity to Intergenerational Social Weights

Benefit rule	Marginal financing	All alive		All gen., 1%		All gen., 2%		All gen., 3%		Future, 2%		All-gen. rank range	
		CEV	Rank	CEV	Rank	CEV	Rank	CEV	Rank	CEV	Rank	Best	Worst
Immediate benefit cut	Funded: capital-income-tax adjustment	0.737	1	0.508	1	0.817	1	0.882	1	0.936	5	1	1
Gradual benefit cut	Funded: capital-income-tax adjustment	0.563	2	0.451	2	0.710	2	0.754	2	0.929	6	2	2
Future-retiree-only cut	Funded: capital-income-tax adjustment	0.357	3	0.366	6	0.523	3	0.519	3	0.768	9	3	6
Gradual benefit cut	Funded: payroll-tax adjustment	0.198	5	0.378	4	0.516	4	0.499	4	0.990	3	3	4
Immediate benefit cut	Funded: payroll-tax adjustment	0.165	6	0.371	5	0.507	5	0.490	5	1.017	2	5	6
Future-retiree-only cut	Funded: payroll-tax adjustment	0.309	4	0.389	3	0.480	6	0.438	7	0.734	10	3	8
Immediate benefit cut	Unfunded: no tax adjustment	0.088	7	0.256	7	0.430	7	0.459	6	0.941	4	5	7
Gradual benefit cut	Unfunded: no tax adjustment	0.021	8	0.224	8	0.379	8	0.401	8	0.913	7	7	8
Future-retiree-only cut	Unfunded: no tax adjustment	-0.283	11	0.093	11	0.156	9	0.149	9	0.811	8	9	11
No reform	Funded: capital-income-tax adjustment	-0.226	10	0.124	10	0.136	10	0.093	10	0.677	11	9	10
No reform	Unfunded: no tax adjustment	-1.108	12	0.132	9	0.075	11	-0.048	12	1.850	1	9	12
No reform	Funded: payroll-tax adjustment	0.000	9	0.000	12	0.000	12	0.000	11	0.000	12	11	12

Notes: CEVs use the fully-funded regime with a payroll tax under no reform as the benchmark. All-generation and future columns, as usual, apply equal weight per person and discount future cohorts. The "All-gen." columns include everyone alive at the announcement date and future worker-entrants. The "Future" column only assesses workers entering the economy after the announcement is made (it also applies the standard 2% discount factor). The final columns give each package's best and worst all-generation rank over annual social discount rates from 0.75 to 5 percent.

The table shows that, in general, immediate and gradual benefit cuts accompanied by capital-income-tax relief are the preferred packages. This remains robust to the range of weights I have allowed to be placed on the welfare of future generations.

B.4 Extrema of Policy Proposal Effects

Table 8: "Winners" and "Losers" by Reform + Financing Package

Benefit Reform	Financing Package	Min. CEV	Biggest Loser	Max. CEV	Biggest Winner	All Gain?
Immediate benefit cut	Funded: payroll-tax adjustment	-1.775	Average incumbent retiree	1.947	Entrant at year 6.75	No
Gradual benefit cut	Funded: payroll-tax adjustment	-1.136	Average incumbent retiree	1.924	Entrant at year 9.25	No
Future-retiree-only cut	Funded: payroll-tax adjustment	-0.784	Entrant at year 0.0	1.636	Average incumbent retiree	No
Immediate benefit cut	Funded: capital-income-tax adjustment	-0.966	Average incumbent retiree	1.905	Average incumbent worker	No
Gradual benefit cut	Funded: capital-income-tax adjustment	-0.531	Average incumbent retiree	1.433	Average incumbent worker	No
Future-retiree-only cut	Funded: capital-income-tax adjustment	-0.927	Entrant at year 0.0	1.918	Average incumbent retiree	No
Immediate benefit cut	Unfunded: no tax adjustment	-2.870	Entrant at year 0.0	1.475	Average incumbent worker	No
Gradual benefit cut	Unfunded: no tax adjustment	-2.774	Entrant at year 0.0	1.198	Average incumbent worker	No
Future-retiree-only cut	Unfunded: no tax adjustment	-2.694	Entrant at year 0.0	2.872	Average incumbent retiree	No

Notes: The relevant groups for comparison are average incumbent workers, average incumbent retirees, and quarterly worker-entry cohorts through the terminal year 40. Every CEV uses the same financing scheme under no reform at the same date. "All gain' asks whether every group analyzed gains over the course of their entire lifetime, as measured by CEV, from the policy package. "Biggest Loser" shows who, as measured by CEV, who was most adversely affected by the policy change. "Biggest Winner" shows who was most favorably affected by the policy change.

No individual financing scheme – reform combination creates an scenario where every group in the model benefits. Each package creates winners, who generate positive CEV over their lifetime, and losers, who end up with negative CEV. Immediate and gradual benefit cuts generally transfer welfare from incumbent retirees toward workers. Future-retiree-only cuts protect incumbent retirees by shifting losses toward workers entering near the announcement. The unfunded implementation gives rise to the strongest intergenerational conflict among the available options as they tend to benefit groups alive at the announcement and simultaneously impose especially large losses on new entrants.