Optimal Income Redistribution

Árpád Ábrahám, Pavel Brendler,
Eva Cárceles-Poveda

Discussion Paper 25/820

15th December 2025

School of Economics

University of Bristol
Priory Road Complex
Bristol
BS8 1TU
United Kingdom

University of BRISTOL
School of Economics

Optimal Income Redistribution*

Árpád Ábrahám
University of Bristol

Pavel Brendler
University of Nottingham

Eva Cárceles-Poveda
Stony Brook University

December 15, 2025

Abstract

We study whether redesigning the social security and income tax-transfer systems can deliver significant welfare gains. Our rich quantitative model features both realistic inequality over the lifecycle and the key channels through which redistributive policies can distort aggregate allocations. We find that there are two distinct ways to achieve significant welfare gains. The first prioritizes reducing distortions through a regressive pension system, resulting in higher inequality. The second reduces inequality through progressive pensions, complemented with a less progressive tax system to mitigate the rise in distortions. In both reform types, pension progressivity emerges as a powerful instrument to either manage distortions or redistribute income within generations. Since redistributive instruments turn out to be highly distortionary in our benchmark economy, the policy reducing distortions turns out to be optimal under utilitarian social welfare.

Keywords: Optimal Progressive Taxation and Pensions; Social Security; Inequality.

JEL codes: E2, E6, H2, H3.

*We are thankful for valuable comments and discussions to Cristina Arellano, Indraneel Chakraborty, Pawel Doligalski, Hans Fehr, Thomas Hintermaier, Juan Carlos Conesa, Jonathan Heathcote, Gueorgui Kambourov, Karen Kopecky, Dirk Krueger, Andre Luduvice, Alexander Ludwig, Chris Phelan, Tom Phelan, Baxter Robinson, Michael Stimmelmayr, Gustavo Ventura, Nicolas Werquin, as well as the seminar participants at numerous seminars and conferences. We thank Oliver Freyermuth from the Physics Department at the University of Bonn for his invaluable IT support. All errors are ours.

1 Introduction

Progressive tax-transfer systems are key instruments used to provide insurance and redistribution. In turn, pay-as-you-go social security systems can also be designed to redistribute income and provide insurance against shocks over the lifecycle. Given the challenges presented by increasing (lifetime) income inequality, the redesign of these two systems is on the agenda in many developed countries. The main contribution of this paper is to build a quantitative framework for the evaluation of such reforms. We show that the welfare benefits of a joint design can be much larger than reforming the income tax-transfer system alone. Moreover, there are two very distinct ways to achieve significant welfare gains: one that prioritizes reducing distortions and another one that focuses on reducing inequality, with pension progressivity emerging as a very powerful instrument to either reduce distortions or redistribute income within cohorts.

To adequately capture the tradeoff between efficiency and redistribution, our framework features the key margins that progressive tax and pension systems could potentially distort. We allow for both endogenous labor supply and endogenous retirement. Workers accumulate human capital through learning-by-doing, implying that current labor supply decisions have long-term consequences on both earnings and pensions. Pensions, which feature an earnings-dependent replacement rate, distort labor supply both through the linear payroll taxes required to finance them and through the replacement rate, which acts as an implicit tax on current labor supply. Asset accumulation is distorted by the two systems, since progressive income taxes reduce the return on savings for wealthier households, while a more generous pension system makes savings less desirable.

With these features, we can address the following pivotal question: To what extent do the income tax-transfer program and the pension system complement each other in delivering the optimal balance between redistribution and incentives? To answer this question, we calibrate our quantitative framework to account for the current income tax and pension systems in the U.S. Specifically, we approximate the income tax schedule in the data with a flexible parametric form following Benabou (2000) and Heathcote et al. (2017), with two parameters that control the income tax progressivity and the average income tax level. As in the U.S. tax system, we assume that the household's taxable income is comprised of both capital income and either labor income for workers or pension benefits for retirees. As for the social security system, we approximate the replacement rate schedule that maps the agent's average lifetime earnings during the working stage to the pension benefit using a similar parametric form, with capped social security earnings (and hence benefits). The remaining model parameters are calibrated to the recent U.S. data to deliver empirically plausible age profiles and distributions of earnings, wealth, and bequests.

In a first set of policy experiments, we allow the Ramsey government to choose permanently the curvature (progressivity) of the replacement rate schedule and the income tax progressivity, either jointly or in isolation. This allows us to compare the results to the existing literature and to understand the distinct redistributive and distortionary properties of the two instruments. In a second set of experiments, we also allow the government to choose the level of the replacement rate schedule (generosity). To evaluate the welfare effects of different policies, our benchmark assumes a utilitarian social welfare function that assigns the same weight to all currently alive generations at the time of the policy change, while the welfare calculations take into account the transition from the initial steady state with the status quo policy to the final steady state implied by the new policy. This involves solving optimization problems in high-dimensional spaces with multiple policy choices, including full transition dynamics in this state space. To do this, we develop a novel numerical approach that computes the transition dynamics and value functions only on a reasonably small multi-dimensional grid of policy parameters and then applies spline-based interpolation in the policy space to evaluate welfare at off-grid policies. This method can be applied to a wide range of computationally intensive optimization problems, an additional contribution of the paper.

When the policymaker can optimize over income tax or pension progressivity only, the optimal policy prescribes a considerable reduction in progressivity. This lowers aggregate distortions and generates welfare gains for both the current and future generations. However, the gains are approximately three times higher in the pension progressivity reform. In this case, the new regressive pension system acts as a highly regressive tax on labor supply, creating strong incentives for the most productive agents to work more and increasing labor supply through a better allocation of labor across income types. At the same time, a lower income tax or pension progressivity redistributes resources from low-income to high-income types, increasing inequality within cohorts in spite of the utilitarian welfare function, which favors redistribution towards low-income types. Hence, reducing distortions outweighs the distributional losses in our framework.

Interestingly, when the policy choice is pension progressivity, there is a local maximum with similar welfare gains for the currently alive but drastically different equilibrium and inequality properties. This case prescribes a significantly higher pension progressivity compared to the status quo, with a considerable decrease in inequality and higher aggregate distortions that generate sizable losses for future generations. It turns out that, when the social planner can set both pension and income tax progressivity, the redistributive benefits of higher pension progressivity and the efficiency benefits of lower income tax progressivity can be combined, and we can both reduce distortions and inequality without creating any significant intergenerational redistribution.

Finally, when we allow the policymaker to optimize also over pension generosity, the key new finding is that pension generosity increases significantly, leading to a large redistribution towards current generations. In fact, both the lower lifecycle savings due to higher pensions and the reduced labor supply incentives due to higher social security taxes contribute to a poorer economy in the long run. As before, the optimal (utilitarian) policy still prescribes a significantly lower pension progressivity, and it achieves large welfare gains and overwhelming political support by considerably mitigating aggregate distortions while creating higher inequality. Moreover, similarly to the combined pension and income tax progressivity reform, the policymaker can also generate very high welfare gains and strong political support from current generations with an alternative policy that increases pension progressivity and reduces within-cohort inequality at the expense of higher aggregate distortions, which are mitigated with a lower income tax progressivity.

To sum up, when the social planner can jointly redesign both the tax-transfer and the pension systems, pension progressivity becomes a powerful instrument for either managing distortions or redistributing within generations. In fact, there are two distinct ways to improve social welfare. The first focuses on reducing distortions and features regressive pensions, resulting in higher inequality (in consumption/welfare). The second focuses on reducing inequality and, therefore, features very progressive pensions complemented with a much less progressive income tax system. While both reforms deliver large welfare gains and have overwhelming political support among the currently alive generations, the winners (and losers) are coming from opposite ends of the income distribution.

To check the robustness of our main findings, we analyze some additional reforms. First, since the full reform generates substantial intergenerational redistribution, we compute the optimal policy both when we limit this type of redistribution and when we redistribute in the opposite direction. In the former case, we impose additional constraints that ensure that no cohort in the present or future is worse off (on average) after the policy change. In the latter case, we use an alternative social welfare function that assigns all the weight to future generations in the final steady state. Our main result regarding the fact that the social planner can deliver significant welfare gains with two distinct types of reforms remains robust. In addition, the optimal policy is the one with significantly lower pension progressivity and more inequality. The only obvious difference is that pension generosity is now increasing much less, since we limit intergenerational redistribution, and it actually decreases below the status quo when we maximize the welfare of future generations.

Finally, to better understand the role of distortions, we also consider (counterfactual) versions of our environment in which some of the channels creating these distortions are eliminated. In all counterfactuals, both the reform that reduces distortions but increases inequality with less progressive pensions and the one with more progressive pensions that reduces inequality still deliver significant welfare gains. The only important difference is that, in cases with significantly less distortions (exogenous human capital accumulation, separate taxation of capital and labor income, exogenous retirement, or partial equilibrium), the redistribution motive dominates, and the optimal policy becomes the one that reduces inequality.

Literature review. Our paper bridges several growing strands of the macroeconomic literature. First, there is a large quantitative macroeconomics literature that analyzes the optimal progressivity of tax-transfer systems, with key contributions by Heathcote et al. (2017), Conesa et al. (2009), Heathcote et al. (2020), Heathcote and Tsujiyama (2021), Carroll et al. (2023), Guner et al. (2023), while Macnamara et al. (2022) adds to the analysis the optimal consumption taxes. This literature convincingly demonstrates that income tax-transfer systems serve as effective tools for insurance and redistribution. We build on these insights and show that allowing for a joint reform of income taxes and pensions can generate substantially larger welfare gains than reforming the tax-transfer system alone.

Second, there exists some research studying the optimal pension progressivity in quantitative lifecycle models following Huggett (1996). Some important examples include Huggett and Parra (2010), Nishiyama and Smetters (2008), Fehr and Habermann (2008), Fehr et al. (2013), Kindermann and Pueschel (2021), and Nam (2024). These studies adopt simpler environments with fewer distortionary channels, focus primarily on the welfare of future generations, and typically examine pension progressivity in isolation from other policy instruments. While most of these papers conclude that more progressive pensions are optimal, Huggett and Parra (2010) and Nishiyama and Smetters (2008) find that lowering pension progressivity may be desirable, as the insurance and redistribution losses are offset by the gains of reducing labor distortions. Our model incorporates multiple sources of distortions and shows that sizable welfare gains can arise from these two types of reforms, with the level of distortions determining whether the one with a lower or higher pension progressivity maximizes social welfare.[1]

Finally, a small but growing literature explores joint reforms of pensions and income taxation. Recent examples include Makarski et al. (2022) and Tran and Zakariyya (2025). These papers differ from ours in several respects: they use a different definition of pen-

[1] In a related paper, Brendler (2020) optimally chooses the generosity of the pension system using Pareto weights that are a function of the household's age.

sion progressivity, do not optimize over pension generosity, and abstract from several key distortionary channels. Brendler (2023) analyzes optimal joint pension generosity and income tax policies in the U.S., focussing on the impact of earnings inequality and demographic transition on the optimal policy mix, but does not consider changes in pension progressivity. These papers generally find that redistribution through pensions can substitute for redistribution through the tax system. We show that this is not generally the case, as our optimal reforms reducing pension progressivity (and distortions) do not require any significant change in tax progressivity. At the same time, our welfare-enhancing reforms reducing inequality show a similar substitutability as in their work, where the increase in pension progressivity is accompanied by a reduction in income tax progressivity.

The rest of the paper is organized as follows. Section 2 presents the model. Section 3 explains the quantitative experiments, and Section 4 the calibration. Our main findings are discussed in Section 5. Section 6 covers robustness, and Section 7 concludes.

2 Model

This section lays out the quantitative model that we will use for our policy analysis. The main building blocks of the model were chosen keeping three important aspects in mind. First, since the paper studies the optimal design of social security (pension) and tax-transfer systems within the institutional framework of the U.S., we model several key features of both systems in relative detail. Second, since a key direction of welfare improvements is the reduction in distortions, our model features several key channels through which both policies distort labor supply and asset accumulation (e.g., human capital accumulation, early retirement, or endogenous bequests). Finally, the other potential source of welfare improvements is improved redistribution or insurance, so we make sure that our model features a reasonable amount of income and wealth inequality, together with a reasonable age profile of income risk.

2.1 Overview

The economy is populated by overlapping generations of agents. Each period, a continuum of agents enter the economy as workers at the age of $j = 1$ (real-life age 25) and continue to work until retirement. Mandatory retirement occurs at age J^R, but agents can retire early at the age of $J^E < J^R$. In case of early retirement, their pension entitlement is reduced. Agents can be either college graduates ($z = z_H$) or non-college graduates ($z = z_L$), and their education remains constant throughout their life. At age 1, agents draw an educational level z from the invariant distribution P_z. Each educational level is associated with an initial amount of skill, denoted by $h_{1,z}$, and an immutable learning ability, denoted by θ_z. At age $j = 1$, agents also draw a fixed effect v that represents unobserved productivity and can also take two values (v_H or v_L). Agents live up to a maximum of $J > J^R$ periods, but they may die earlier due to stochastic mortality. We denote by $\psi_j^{z,v}$ the probability that a type (z, v) agent survives up to age $j + 1$, conditional on surviving up to age j. The population grows at an exogenous rate of n, with the population size normalized to 1 in the initial period ($t = 0$). All aggregate variables, prices, and policy variables will be indexed by time, since our analysis considers deterministic transitional dynamics. Also, aggregate variables at t (except prices) are divided by $(1 + n)^t$ to express them in per capita terms.

Financial markets are incomplete, as there are no insurance markets against idiosyncratic mortality and labor productivity shocks. Agents can only self-insure against shocks by accumulating shares in a representative firm and government bonds. Since both assets bear no risk and generate the same return, agents are indifferent between them, and their total asset wealth is denoted by a. Borrowing is ruled out, and agents enter the economy with no assets unless they receive bequests from their parents, as explained below.

2.2 Households

Households have time-separable preferences and a constant discount factor β. They are endowed with one unit of productive time each period, which they can allocate between work l and leisure $1 - l$. The utility from consumption and leisure in each period is given by the function $u(c, 1 - l)$. Agents also derive utility $q(a')$ directly ("warm glow") from leaving a bequest of a' in case of death.

For an agent that works l hours, the pre-tax earnings are given by

$$e = w_{z,t} h_{j,z} v y_j l, \tag{1}$$

where the deterministic part of the hourly wage $w_{z,t} h_{j,z} v$ is governed by the wage rate $w_{z,t}$ (determined in equilibrium), the skill level $h_{j,z}$ and the fixed effect v, while the stochastic part is given by an idiosyncratic shock y_j that varies over the lifecycle.

As for the evolution of the skill level, we assume that human capital is accumulated endogenously through the following learning-by-doing mechanism:

$$h_{j+1,z} = (1 - \delta_h) h_{j,z} + \theta_z (h_{j,z} l)^{\gamma_h}, \tag{2}$$

which implies that effective labor supply in the current period, $h_{j,z}l$, contributes to future productivity. The parameters θ_z (learning ability) and $\gamma_h \in (0,1)$ are related to the speed at which the agent builds up skills, and δ_h is the rate at which skills depreciate. This learning-by-doing human capital accumulation formulation brings a strong dynamic component to labor supply decisions, making any direct or indirect (through pension progressivity) form of labor income taxation more distortionary.[2]

Finally, the stochastic part of the worker's productivity is driven by the idiosyncratic productivity shock y_j, which follows an AR(1) process:

$$\log y_j \quad = \quad \rho_y \log y_{j-1} + \epsilon_j \text{ with } y_1 = 1 \text{ and } \epsilon_j \sim \mathcal{N}(0, \sigma_\epsilon^2). \tag{3}$$

The presence of stochastic shocks will generate an insurance motive for government policy, while the fixed agent characteristics (z and v) will generate a redistributive motive.

2.3 Social security and income tax-transfer system

Before describing the household's decision problem, this subsection discusses how social security and the income tax-transfer system are introduced into the household's problem. The main normative exercise of the paper will be to redesign these policy tools to find an optimal balance between (production) efficiency and redistribution/insurance. When doing this, we will assume that the government has two separate budget constraints. Social security taxes balance the social security budget every period, while the other taxes and new debt issuance are used to finance government spending and debt repayments.

2.3.1 Social security

The specific features of the social security system closely follow the key aspects of the U.S. social security system. All workers pay a proportional social security tax $\tau_{SS,t}$ on their *taxable earnings*[3], which are given by $\min(e, cap)$, where cap is an exogenous maximum taxable earnings threshold and e are the total pre-tax earnings defined in (1).

Workers must retire once they reach the mandatory retirement age J^R, but they already become eligible for pension benefits at the age of $J^E < J^R$ and, therefore, may choose to retire early. Since the timing of retirement will affect the benefit amount, we introduce

[2] Holter et al. (2019) already shows the importance of the endogenous human capital channel in the context of the progressivity of the income tax system.

[3] We refer to *"taxable earnings"* as the base of social security taxes and of pension entitlements. In contrast, we will refer to the tax base of the tax-transfer system as *"taxable income"* (see the definition in (7) below).

an individual state variable j^R that is equal to the age at which the agent has retired. For workers, this variable is set to $j^R = 0$.

Average taxable lifetime earnings during the working stage, $\bar{e}_j$, are equal to the simple average of past taxable earnings, and they can be computed recursively as follows:

$$\bar{e}_{j+1} = \left[(j-1)\bar{e}_j + \min(e_j, cap)\right]/j. \tag{4}$$

All workers enter the model with no prior earnings histories, i.e., $\bar{e}_1 = 0$, and once the agent retires, $\bar{e}_j$ remains constant. During retirement, the agent receives a pension benefit $p = p(\bar{p}_t, j^R)$ which depends on two components: i) the full pension amount $\bar{p}_t$ for which the agent would have qualified had they retired at the mandatory retirement age J^R and ii) the actual retirement age j^R.

The full pension amount is given by $\bar{p}_t(\bar{e}; \boldsymbol{\lambda}_{SS,t}) = \bar{e} \times R_t(\cdot)$, where $R_t(\cdot)$ is the replacement rate schedule. The statutory schedule of the average replacement rate in the data as a function of average lifetime earnings $\bar{e}$ is approximated by a simple parametric class:

$$R_t(\bar{e}, j^R; \boldsymbol{\lambda}_{SS,t}) = \begin{cases} \bar{\lambda}_{SS,t} \times \tilde{e}^{1-\lambda_{SS,t}} & \text{if } \tilde{e} \geq \tilde{e}_{\min} \\ \bar{\lambda}_{SS,t} \times \tilde{e}_{\min}^{1-\lambda_{SS,t}} & \text{otherwise,} \end{cases} \tag{5}$$

where $\tilde{e} = \bar{e}/\mathcal{E}_{t-j+j^R}$, $\mathcal{E}_{t-j+j^R}$ are the average (taxable) lifetime earnings among retirees at the time when the agent enters retirement, $\tilde{e}_{\min}$ is the first bracket of the statutory replacement rate schedule and $\boldsymbol{\lambda}_{SS,t} = \{\lambda_{SS,t}, \bar{\lambda}_{SS,t}\}$ are the two policy variables associated with pensions. In particular, $\bar{\lambda}_{SS,t}$ controls the level of the replacement rate schedule or pension generosity and $\lambda_{SS,t}$ governs the progressivity of the pension system.[4]

Figure 1 shows how the statutory schedule of average replacement rates depends on the pension generosity $\bar{\lambda}_{SS,t}$ (left panel) and the pension progressivity (middle and right panels). When $\lambda_{SS,t} > 1$, replacement rates decrease in the worker's (normalized) average lifetime earnings, and the pension system is referred to as progressive (central panel). When $\lambda_{SS,t} < 1$, the opposite is the case, and the pension system is regressive (right panel).[5]

When the agent retires early, the pension payment is reduced proportionately to the number of years prior to the mandatory retirement age. The penalty function for early

[4] The individual's average lifetime earnings in the replacement rate schedule are normalized by the average taxable earnings among retirees at the time when the agent enters retirement, $\mathcal{E}_{t-j+j^R}$. Hence, $\bar{\lambda}_{SS,t}$ corresponds to the replacement rate of an agent whose average lifetime earnings, at the time of retirement, are equal to the average (past) taxable earnings of all retirees in that period.

[5] Another element of pension progressivity of the U.S. system is captured by the cap. Households are not entitled to receive any pension benefits beyond the cap, but the social security tax is not collected above the cap either. For the rest of the analysis, we keep the cap constant.

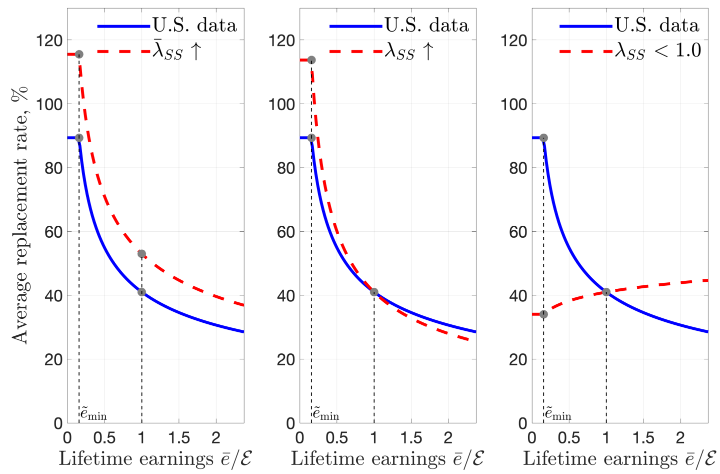

Figure 1: Statutory schedule of average replacement rates.

Notes: The solid lines show the schedule calibrated to the current U.S. data ($\bar{\lambda}_{SS} = 0.44$ and $\lambda_{SS} = 1.42$). The dashed line in the left panel shows the schedule for a higher generosity ($\bar{\lambda}_{SS} = 0.53$). The dashed line in the central panel shows the schedule for higher progressivity ($\lambda_{SS} = 1.55$), and the right panel depicts a regressive schedule ($\lambda_{SS} = 0.9$). The lowest (normalized) bend point, $\tilde{e}_{\text{min}} = \bar{e}_{\text{min}}/\mathcal{E}$, is equal to 0.11.

retirement is specified as:

$$p(\bar{p}_t, j^R) = (1-\delta_p)\bar{p}_t + \left(\frac{j^R - J^E}{J^R - J^E}\right)\delta_p\bar{p}_t. \tag{6}$$

If retirement occurs at the earliest age ($j^R = J^E$), the agent loses (permanently) a fraction δ_p of the full pension amount. If they retire later but before the mandatory age ($J^R > j^R > J^E$), the loss is scaled down proportionally. When the agent retires at the mandatory retirement age ($j^R = J^R$), the expression above boils down to the full pension amount.[6]

2.3.2 Income tax-transfer program

In addition to a consumption tax τ_c, households face a non-linear income tax-transfer scheme. The agent's *taxable income* is given by

$$\iota = \begin{cases} r_t a + e - 0.5\tau_{SS,t}\min(cap, e) & \text{if } j^R = 0 \\ r_t a + p(\bar{p}_t(\bar{e}; \boldsymbol{\lambda}_{SS,t}), j^R) & \text{otherwise.} \end{cases} \tag{7}$$

The first line in the previous equation is the taxable income of workers, where e are pre-tax earnings defined in (1) and $r_t a$ is the return on savings. Since the part of labor income

[6]The specification approximates reasonably well the penalty in the current social security legislation.

that is paid by the employee as social security contributions is not subject to income taxes, a corresponding portion is deducted from the worker's taxable income. The second line of the expression above corresponds to the taxable income of retired agents. In line with the data, pension benefits are subject to income taxation.[7]

Income is taxed according to the income tax function $\Psi(\cdot)$. Following Benabou (2000) and Heathcote et al. (2017), the net tax liability is specified as

$$\Psi_t(\iota; \boldsymbol{\lambda}_{I,t}) = \iota - \mathcal{I}_t \cdot (1 - \bar{\lambda}_{I,t}) \cdot (\iota/\mathcal{I}_t)^{1-\lambda_{I,t}}, \tag{8}$$

where ι is the agent's taxable income defined in (7), $\mathcal{I}_t$ is the economy-wide average taxable income, and $\boldsymbol{\lambda}_{I,t} = (\lambda_{I,t}, \bar{\lambda}_{I,t})$ are the policy variables associated with the income tax-transfer program. The agent's income after taxes and transfers is then equal to $\iota - \Psi_t(\iota; \boldsymbol{\lambda}_{I,t}) = \mathcal{I}_t \cdot (1 - \bar{\lambda}_{I,t}) \cdot (\iota/\mathcal{I}_t)^{1-\lambda_{I,t}}$.

As with pensions, $\bar{\lambda}_{I,t}$ controls the average income tax level, while $\lambda_{I,t}$ determines income tax progressivity. When $0 < \lambda_{I,t} < 1$, the income tax system is progressive, in the sense that the marginal tax increases in income and the average tax rate is higher (lower) than $\bar{\lambda}_{I,t}$ for any agent whose income is above (below) $\mathcal{I}_t$. Whenever $\lambda_{I,t} > 0$, the tax liability $\Psi_t(\iota; \boldsymbol{\lambda}_{I,t})$ can be negative for low levels of income, implying that households with low income are entitled to transfers. This feature reflects the negative income tax many low-income households face through the Earned Income Tax Credit (EITC) program.

As in the U.S. tax code, our income tax specification does not discriminate between different income sources (labor income, capital income, and pension benefits). In addition, the only (major) sources of income for retirees are pension and capital (interest) income. These sources are not subject to the EITC that targets labor income. Given this, we adjust the tax liabilities for retirees as follows:

$$\Psi_t^R(\iota; \boldsymbol{\lambda}_{I,t}) = \max\{0, \iota - \mathcal{I}_t \cdot (1 - \bar{\lambda}_{I,t}) \cdot (\iota/\mathcal{I}_t)^{1-\lambda_{I,t}}\}. \tag{9}$$

This restriction implies that retirees with low taxable income will not pay income taxes, but they are also not entitled to transfers.[8]

[7]Note that, in the closely related literature, Heathcote et al. (2017) and Heathcote et al. (2020) abstract away from capital income, while Conesa et al. (2009) and Krueger and Ludwig (2016) assume that capital income is taxed separately. All of the mentioned studies do not tax pension benefits.

[8]This assumption allows us to use the same parametrization of the tax function for both workers and retirees and, hence, keep the set of policy variables in the optimal policy analysis tractable.

2.4 Household's problem

Below, we spell out the agent's optimization problems at different stages of the lifecycle.

Problem of workers of age $1, ..., J^E - 2$

Workers become eligible for early retirement benefits once they reach the age of J^E. The retirement decision must be made one period in advance, which implies that agents do not make any retirement decision during the first $J^E - 2$ periods. During this stage, workers choose how much to consume, how much to save, and how to split the unitary time endowment between leisure and work. In each period, agents draw the labor productivity shock y before making any decisions.

We denote by $\boldsymbol{x} = (j, j^R, z, v, y, h, \bar{e}, a)$ the agent's current state, where j is age, j^R – retirement age, z – educational level, v – fixed effect, y – idiosyncratic AR(1) productivity shock, h – human capital, $\bar{e}$ – average lifetime earnings, and a – assets. Since the agent remains in the labor force next period, their retirement status next period will remain unchanged: $j^{R\prime} = j^R = 0$. Moreover, given the optimal choice of hours worked l, the agent's next period stock of human capital h' and average lifetime earnings $\bar{e}'$ will follow from the law of motions (2) and (4), respectively. Note that a, the asset position of the household at the beginning of the period, is composed of both the savings of the last period and the bequests b received at the end of the last period.

Let $V_t(\boldsymbol{x})$ denote the agent's value function at time t and state $\boldsymbol{x}$. The agent's optimization problem is given by

$$\begin{aligned} V_t(\boldsymbol{x}) &= \max_{\substack{c, a' \geq 0, \\ l \in [0,1]}} \Bigg\{ u(c, 1-l) + \beta \psi_j^{z,v} \mathbb{E}_{y'|y} \Bigg[\Gamma_{j+1,t+1}^{z,v} \int V_{t+1}(\boldsymbol{x}'_b) \Phi_{t+1}^{z,v}(b) db + \\ & \qquad + (1 - \Gamma_{j+1,t+1}^{z,v}) V_{t+1}(\boldsymbol{x}') \Bigg] + (1 - \psi_j^{z,v}) q(a') \Bigg\} \end{aligned} \tag{10}$$

subject to the budget constraint

$$a' + (1 + \tau_c) c = (1 + r_t) a + e - \tau_{SS,t} \min\{e, cap\} - \Psi_t(\iota; \boldsymbol{\lambda}_{I,t}). \tag{11}$$

In the budget constraint, e are pre-tax earnings defined in (1), ι is taxable income defined in (7), and $\Psi_t(\cdot)$ is the income tax liability defined in (8). An agent of type (z, v) and age j dies with probability $1 - \psi_j^{z,v}$, in which case the agent derives utility from leaving a bequest of a' to their "offspring", represented by the last term $q(a')$ in the problem above. Given this, bequests are left both accidentally (when households die earlier than expected) and intentionally due to this "warm glow" motive. Also, some agents will not leave any bequests because they live longer than expected and need to use the remaining assets for their own consumption or because they do not find it optimal to leave any bequest.

When modeling the distribution of bequests among surviving agents, we build upon Nishiyama and Smetters (2007), who assume that bequests are distributed randomly. Their approach accounts for the empirical fact that many households do not receive any bequests during their lifetime, while others receive multiple bequests. We capture this in (10) by assuming that, if the agent survives, which occurs with probability $\psi_j^{z,v}$, there is a probability $\Gamma(\cdot)$ that they will receive a bequest next period and a probability $1-\Gamma(\cdot)$ that they will receive none. Given this, $V_{t+1}(\boldsymbol{x}')$ is the value function of the agent at time $t+1$ if they do not receive any bequests, and $V_{t+1}(\boldsymbol{x}'_b)$ is the value function of the agent at $t+1$ if they receive a bequest b, with $\boldsymbol{x}' = (j+1, j^{R\prime}, z, v, y', h', \bar{e}', a')$ and $\boldsymbol{x}'_b = (j+1, j^{R\prime}, z, v, y', h', \bar{e}', a'+b)$. We extend their approach in two ways to account for additional empirical facts. First, we assume type and age dependence in the probability of receiving bequests, i.e., $\Gamma^{z,v}_{j+1,t+1}$ varies with type (z,v) and j. This replicates the fact that higher-income households are more likely to receive bequests than lower-income ones, and it ensures that the age profile of bequest receipts aligns with the data. Second, we account for the fact that income-richer households receive larger bequests, which is captured by the type-specific distribution of received bequests, $\Phi^{z,v}_{t+1}$. Of course, both $\Gamma^{z,v}_{j+1,t+1}$ and $\Phi^{z,v}_{t+1}$ are endogenous and, in equilibrium, they are determined by the lifecycle savings behavior of households who leave these bequests (see Appendix A.1 for more details).

Problem of workers of age $J^E-1, ..., J^R-1$

Once an agent reaches the age of J^E-1, at the end of the period, they may choose to enter retirement next period before observing their next period productivity shock, and we assume that retirement is an absorbing state. In this case, welfare is given by

$$
\begin{aligned}
V_t(\boldsymbol{x}) \quad = \quad & \max_{\substack{c,a'\geq 0,\\ l\in[0,1]}} \Bigg\{ u(c,1-l) + \beta\psi_j^{z,v}\left[\Gamma^{z,v}_{j+1,t+1}\int \tilde{V}_{t+1}(\boldsymbol{x}'_b)\Phi^{z,v}_{t+1}(b)db + (1-\Gamma^{z,v}_{j+1,t+1})\tilde{V}_{t+1}(\boldsymbol{x}')\right] + \\
& +(1-\psi_j^{z,v})q(a')\Bigg\}, \qquad (12)
\end{aligned}
$$

with

$$
\tilde{V}_{t+1}(\cdot) = \max\{V^R_{t+1}(\cdot), \mathbb{E}_{y'|y}\left[V_{t+1}(\cdot)\right]\} \qquad (13)
$$

and subject to the budget constraint in(11). In the previous equation, $V^R_{t+1}(\cdot)$ is the value function of the agent if they retire next period with the appropriate future state vector and $j^{R\prime} = j+1$ (see (14) below). At age $J^R - 1$, the agent will retire with certainty next period and $\tilde{V}_{t+1}(\boldsymbol{x}')$ and $\tilde{V}_{t+1}(\boldsymbol{x}'_b)$ in (12) will be replaced by $V^R_{t+1}(\boldsymbol{x}')$ and $V^R_{t+1}(\boldsymbol{x}'_b)$, respectively. Note also that the stock of human capital h and the productivity shock y do not affect household decisions and welfare during retirement, while their type (z, v) affects the survival probability rate $\psi^{z,v}_j$ and the probability of receiving bequests $\Gamma^{z,v}_{j+1,t+1}$.

Problem of retirees (age $J^E, ..., J$)

Retirees devote their unitary time endowment to leisure. Their maximization problem is

$$\begin{aligned} V^R_t(\boldsymbol{x}) \quad = \quad & \max_{c,a'\geq 0}\Bigg\{ u(c,1) + \beta\psi^{z,v}_j \left[\Gamma^{z,v}_{j+1,t+1} \int V^R_{t+1}(\boldsymbol{x}'_b)\Phi^{z,v}_{t+1}(b)db + (1-\Gamma^{z,v}_{j+1,t+1})V^R_{t+1}(\boldsymbol{x}') \right] \\ & + (1-\psi^{z,v}_j) q(a') \Bigg\} \end{aligned} \tag{14}$$

subject to the budget constraint

$$a' + (1+\tau_c)c = (1+r_t)a + p(\bar{p}_t(\bar{e};\boldsymbol{\lambda}_{SS,t}), j^R) - \Psi^R_t(\iota;\boldsymbol{\lambda}_{I,t}). \tag{15}$$

For retirees, both the retirement age indicator and the lifetime earnings remain unchanged, namely $j^{R\prime} = j^R$ and $\bar{e}' = \bar{e}$.

2.5 Production technology

The economy is populated by a firm that rents capital and labor from households. We assume that college graduate workers and non-college graduate workers are imperfect substitutes in the production function, but workers of the same education are perfect substitutes across different ages. Let $N_{z,t}$ denote the aggregate labor of type z at time t measured in efficiency units. Then, the total effective labor at time t is given by

$$N_t = \left(N^{\rho}_{L,t} + N^{\rho}_{H,t}\right)^{\frac{1}{\rho}}, \tag{16}$$

where $1/(1-\rho)$ is the elasticity of substitution between the two worker types.

The representative firm produces the final output good according to

$$Y_t = Z K^{\varpi}_t N_t^{1-\varpi}, \tag{17}$$

where Z is the total factor productivity (TFP), K_t is the aggregate capital stock, and $\varpi \in (0,1)$ measures the elasticity of output with respect to the input of capital services. The output Y_t can be consumed or invested in capital. The firm rents capital and hires labor from households on competitive spot markets. Capital depreciates at rate δ. This implies that the interest rate r_t and the wage rates (skill prices) $w_{z,t}$ are determined by the firm's profit maximization conditions:

$$
\begin{aligned}
w_{z,t} &= (1-\varpi) Z \left(K_t/N_t\right)^{\varpi} \left(N_t/N_{z,t}\right)^{1-\rho} \text{ for } z \in \{H, L\}, && (18) \\
r_t &= \varpi Z \left(K_t/N_t\right)^{\varpi-1} - \delta. && (19)
\end{aligned}
$$

As long as $\rho < 1$, both types of workers are imperfect substitutes and the first condition implies that relative skill prices in equilibrium, $w_{H,t}/w_{L,t} = (N_{L,t}/N_{H,t})^{1-\rho}$, co-move with their relative supply.

2.6 Equilibrium

The equilibrium of this economy is defined in detail in Appendix A.2. The equilibrium is consistent with the optimal decisions of households and firms, given equilibrium prices, bequests, and government policies. It also guarantees that, given prices and optimal decisions, markets clear and the distribution of bequests left is consistent with the distribution of bequests received. For these conditions, a key object is $F_{t,j}(\boldsymbol{x})$ defined as the distribution of households across the state $\boldsymbol{x}$ for a given age j at time t.

We impose budget clearing for the two government systems, modelling the social security budget independently of the federal budget, in line with the U.S. legislation. We assume that the social security tax, $\tau_{SS,t}$, balances the social security budget every period:

$$
\tau_{SS,t} \sum_j \int_{\boldsymbol{x}:j^R=0} \min(cap, e) dF_{t,j} = \sum_j \int_{\boldsymbol{x}:j^R \neq 0} p(\bar{p}_t(\bar{e}; \boldsymbol{\lambda}_{SS,t}), j^R) dF_{t,j}. \tag{20}
$$

In addition, the government finances an exogenous stream of spending G_t (wasted in the context of our framework) with exogenous debt issuance $D_{t+1} > 0$ and other tax revenue. The income tax level, $\bar{\lambda}_{I,t}$, balances the government budget every period:

$$
G_t + (1+r_t)D_t = \sum_j \int_{\boldsymbol{x}} \Psi_t(\iota; \boldsymbol{\lambda}_{I,t}) dF_{t,j} + \tau_c C_t + (1+n) D_{t+1}, \tag{21}
$$

where C_t is the aggregate consumption level.

3 Quantitative experiments

One of the main objectives of this paper is to assess in a systematic way the social desirability of different combinations of pension and tax-transfer systems. This is a notoriously cumbersome problem, as the characteristics of the "optimal" joint system may depend on what weights we allocate to current and future generations, as well as on other "societal" constraints we might want to impose. In this section, we explain the key elements of the social planner's optimization problem we have adopted.

At time t, the economy is assumed to be in a steady-state equilibrium with a given pension policy $(\bar{\lambda}^0_{SS}, \lambda^0_{SS})$ and income tax progressivity λ^0_I. To simplify notation, we denote the vector of initial policy by $\mathbf{\Lambda}^0 = (\bar{\lambda}^0_{SS}, \lambda^0_{SS}, \lambda^0_I)$. These initial steady-state policy parameters correspond to the status quo and are, therefore, set to their values in the U.S. data. In the same period, the government makes an unanticipated and credible announcement that it will implement joint social security and income tax reforms that will change the status quo policy to a new constant value $\mathbf{\Lambda}^\star$ effective the following period.[9] As explained in the quantitative section later on, we will conduct a series of experiments in which either some or all the policy parameters in $\mathbf{\Lambda}^\star$ will be allowed to change.

When implementing the policy reforms, we make an important assumption that the income tax reform applies to all agents, while the pension system reform only applies to those agents who have not entered retirement at the time of the reform announcement. In addition, the new pension system is phased in linearly over a period of 40 years. This assumption implies that only the new labor market entrants experience the full impact of the pension reform in terms of pension entitlement; for households that contributed to the old system for all or part of their working life, pension benefits will be determined by combining both systems with the appropriate weights.

Clearly, a change in the policy parameters will trigger transitional dynamics for all the aggregates (such as interest rates, wages, or taxes) from the pre-reform to the final steady state associated with the new policy, which we take into account when evaluating the welfare effects of each reform. During this transition, given the new (constant) replacement rate schedule $(\bar{\lambda}_{SS}, \lambda_{SS})$, the social security tax rate $\tau_{SS,t}$ will adjust to balance the social security government budget in (20). Moreover, given the new (constant) income tax progressivity λ_I, the income tax level $\bar{\lambda}_{I,t}$ will balance the government budget in (21).

When implementing the optimal policy, we focus on a social welfare function that max-

[9] An implementation lag is common for social security and tax reforms. It implies that agent decisions in the initial period change only due to anticipation and the resulting general equilibrium (price) effects.

imizes the welfare of currently alive (working and retired) generations:[10]

$$W(\boldsymbol{\Lambda}^0, \boldsymbol{\Lambda}) = \sum_j \int V_0(\boldsymbol{x}; \boldsymbol{\Lambda}^0, \boldsymbol{\Lambda}) dF_{0,j}. \tag{22}$$

In the benchmark case, we consider unconstrained optimization problems in which the planner at time 0 chooses a constant future policy $\boldsymbol{\Lambda}^\star$ that solves:

$$\boldsymbol{\Lambda}^\star = \arg\max_{\boldsymbol{\Lambda}} W(\boldsymbol{\Lambda}^0, \boldsymbol{\Lambda}). \tag{23}$$

Since this assigns no weight to future generations, the optimal policies may deliver considerable redistribution across generations. Given this, in the robustness section, we conduct a second set of policy experiments, which we denote as constrained optimal policies, assuming that the government at time 0 solves the same problem as in (23), but it faces Pareto constraints that limit redistribution. To do this, we require that the optimal policy does not make any cohort, on average, worse off. These intergenerational Pareto constraints can be formalized as follows:[11]

$$\int V_0(\boldsymbol{x}; \boldsymbol{\Lambda}^0, \boldsymbol{\Lambda}^\star) dF_{t,j} \geq \int V_0(\boldsymbol{x}; \boldsymbol{\Lambda}^0, \boldsymbol{\Lambda}^{\mathbf{0}}) dF_{0,j} \tag{24}$$

for all $j = 1, ..., J$, and

$$\int V_s(\boldsymbol{x}^N; \boldsymbol{\Lambda}^0, \boldsymbol{\Lambda}^\star) dF_{s,j=1} \geq \int V_s(\boldsymbol{x}^N; \boldsymbol{\Lambda}^0, \boldsymbol{\Lambda}^{\mathbf{0}}) dF_{s,j=1} \tag{25}$$

for all $s = 1, ..., \infty$, where $\boldsymbol{x}^N$ is the state of a newborn agent. These constraints require that, under the new policy $\boldsymbol{\Lambda}^\star$, no cohorts of the currently living generations (24) and no newborns among future generations (25) are worse off compared to the status quo.[12] Note that the previous two constraints limit the amount of intergenerational redistribu-

[10] We also present the results with a social welfare function that maximizes the welfare of households born into the final steady state (Appendix D.2). We show that most of our key findings are robust to this popular welfare criterion.

[11] Our paper is not the only one trying to design Pareto-improving reforms by imposing these types of constraints. For example, Fehr et al. (2013) and Fehr and Habermann (2008) introduce a lump-sum redistribution authority in the spirit of Auerbach and Kotlikoff (1987) and Nishiyama and Smetters (2007, 2005) to compensate potential losers of reforms by reallocating resources from the winners. In addition, Conesa and Garriga (2008) impose similar constraints to the ones above to study the optimal Pareto-improving design of social security reforms in an environment without idiosyncratic shocks.

[12] Mathematically, these constraints imply positive welfare weights on a future generation whenever the constraint (25) binds for that particular cohort. In fact, the weight is given by the Lagrange multiplier associated with (25) for that cohort. In this sense, these restrictions are similar to other approaches that assign welfare weights exogenously, but our approach provides a clear interpretation of such weights.

tion. However, we may want to limit instead – or in addition – the redistribution within currently living and/or future newborn households by imposing the previous constraints on a subset of education and fixed productivity types (z, v).

There are two important numerical challenges in our problem. First, it has a large state space with three continuous state variables $(a, h, \bar{e})$ and many generations. Second, we need to solve an optimization problem in three policy variables $\mathbf{\Lambda}$, including full transition dynamics in this state space. Moreover, we also impose "Pareto constraints" that involve value functions of multiple generations and types. To tackle these issues, we have developed a numerical approach that requires solving only a limited (albeit not small) number of model versions with different policy parameters, which form a "base" for the optimization, thereby relaxing some of the numerical burden. This approach can be applied to a wide range of computationally intensive optimization problems, and we consider it an additional contribution of the paper. See details in Appendix A.3.

4 Calibration

The model parameters are calibrated to 2010–2019. The economy is assumed to be in a steady-state equilibrium so the time index is dropped throughout this section. The model period equals one year. Agents enter the model as workers at the real-life age of 25. A large set of parameters is calibrated using the Current Population Survey (CPS). An agent in the model corresponds to a household head in the CPS. The agent's educational type z corresponds to the household head's educational level.[13] Table 1 lists the model parameters, dividing them into parameters calibrated outside the model and parameters calibrated inside the model. For the latter parameters, empirical targets are shown in brackets. Below, we explain the calibration strategy in detail.

Demographics. The maximum possible age is $J = 76$ (real-life age of 100). The mandatory retirement age, J^R, is equal to 42 (real-life age of 66), while the early retirement age is $J^E = 38$ (real-life age of 62). The share of college graduates at birth, P_H, is set to 44%, consistent with our CPS sample. The age- and type-specific mortality rates, $1 - \psi_j^{z,v}$, are estimated inside the model to target the age profile of survival probability rates as well as the mortality rates by income at the real-life age of 70 (Appendix B.1). Given $\psi_j^{z,v}$, the birth rate n is set to match the dependency ratio in the data, defined as the ratio of household heads aged 66–100 to household heads aged 25–65, equal to 25%.

Preferences and bequest parameters. The agent's period utility u is assumed to be a

[13]We refer to household heads with a completed college degree or higher as *college graduates* and all remaining households as *high-school graduates*.

Table 1: Model parameters.

Parameter	Interpretation	Value
Parameters determined outside the model:		
Demographics and preferences		
(J, J^E, J^R)	Maximum age, early and normal retirement age	$(76, 38, 42)$
P_H	Share of college graduates at birth	0.44
σ	Coefficient of relative risk aversion	2.0
Labor productivity		
$(\rho_y, \sigma_\epsilon^2)$	Persistence and variance of AR(1) shock	$(0.979, 0.015)$
Production		
(ϖ, δ)	Capital share and capital depreciation	$(0.33, 0.08)$
ρ	Elasticity of substitution is $1/(1-\rho)$	0.285
Government policies		
λ_I	Income tax progressivity	0.216
λ_{SS}	Pension system progressivity	1.42
τ_c	Consumption tax, %	4.1
(d_y, g_y)	Debt-to-GDP ratio and Wasted spending-to-GDP ratio, %	$(100.0, 7.8)$
Parameters calibrated inside the model (target in brackets):		
Demographics		
$\{\psi_j^{z,v}\}$	Type-specific age profile of survival probabilities (mortality rates by income at age 70)	Appendix B.1.
n	Population growth rate, % (dependency ratio = 25%)	1.3
Preferences and bequest parameters		
β	Discount factor $(K/Y = 3.0)$	0.99
γ	Weight on consumption $(\bar{l} = 0.4)$	0.35
ϕ_1	Degree of altruism $(B/A = 1.0\%)$	11.0
ϕ_2	Degree of bequests as luxury good (share of deceased leaving zero bequests = 20%)	4.0
$\Gamma_j^{z,v}$	Type- and age-specific probabilities to receive a bequest	Appendix B.2.
$\Phi^{z,v}$	Type-specific distribution of received bequests	Appendix B.2.
Labor productivity		
γ_h	Elasticity of human capital production (labor supply elasticity = 0.5)	0.4
$(h_{1,H}, h_{1,L})$	Initial skill levels (age profile of earnings from Figure 2)	$(1.2, 0.45)$
(θ_H, θ_L)	Learning ability levels (age profile of earnings from Figure 2)	$(0.17, 0.08)$
δ_h	Skill depreciation, % (age profile of earnings from Figure 2)	5.0
(v_H, v_L)	Fixed effect (Gini index for pre-gov. earnings = 0.40)	$(1.55, 0.64)$
Production		
Z	TFP (average wage rate $w = 1.0$)	0.267
Government policies		
$\bar{\lambda}_{SS}$	Replacement rate level $(\tau_{SS} = 10.6\%)$	0.413
$\bar{e}_{\min}$	Lowest bend point $(\bar{e}_{\min}/\mathcal{E} = 11\%)$	0.05
cap	Maximum taxable earnings (share of taxable earnings above cap = 8%)	1.11
δ_p	Penalty for early retirement (share of retired agents at age 62 = 26%)	0.167

constant relative risk aversion function given by

$$u(c, 1-l) = \frac{\left[c^{\gamma}(1-l)^{1-\gamma}\right]^{1-\sigma} - 1}{1-\sigma}, \tag{26}$$

where σ controls the degree of relative risk aversion and γ is the relative weight on consumption. We fix σ at 2.0 and calibrate γ inside the model to match the fraction of disposable time devoted to work, which is 40.0% in the data.[14] The discount factor β is calibrated internally to match the capital-to-GDP ratio, K/Y, equal to 3.0. The utility from leaving bequests is specified as:

$$q(a') = \phi_1 \frac{(\phi_2 + a')^{1-\sigma}}{1-\sigma}. \tag{27}$$

Here, ϕ_1 is the degree of altruism and ϕ_2 captures how much bequests are a luxury good. We calibrate the two parameters inside the model to match several moments of the bequest distribution in the data. In particular, the degree of altruism ϕ_1 is calibrated to match a bequest-to-wealth ratio, B/A, of 1.0%, which lies in the range reported by De Nardi and Yang (2016). Moreover, we calibrate ϕ_2 to match the share of deceased agents who leave no bequests, equal to 20%. This moment comes from Hurd and Smith (2001), who use the Health and Retirement Study and the Asset and Health Dynamics of the Oldest Old (AHEAD).[15] Finally, we calibrate the age- and type-specific probabilities of receiving a positive bequest, $\Gamma_j^{z,v}$, and the type-specific distribution of received bequests, $\Phi^{z,v}$, inside the model to ensure that the resulting age profile of these probabilities, averaged across types, is consistent with the data. Appendix B.2 explains the details.

Human capital accumulation. We calibrate the parameters of the law of motion for skills in (2) as follows. The speed at which agents acquire skills, $\gamma_h = 0.4$, is calibrated inside the model to match a labor supply elasticity of 0.5, which is in the range of the empirical estimates in the literature on the Frisch elasticity of labor supply, as discussed in Keane (2011). To obtain the model estimate of this elasticity, we follow Badel et al. (2020), who run a regression of the model-generated change in log hours on the change in log wages following MaCurdy (1981).[16] The learning abilities (θ_L, θ_H), the initial stock of skills $(h_{1,L}, h_{1,H})$, and the depreciation rate δ_h are calibrated inside the model to match the profiles of pre-government earnings by age and education in the CPS. To obtain the

[14] Ábrahám and Carceles-Poveda (2008) find in the Time Use Surveys that the household's disposable time is 97 hours per week after deducting sleep and personal care, summing up to 5,096 hours annually, which we assume as the total disposable time when computing the fraction of time devoted to work.

[15] See their Table 11.1, column "All", which reports percentiles of the bequest distribution based on the pooled sample of single and married households.

[16] As the authors, we run IV regressions with a two-stage least squares estimator, using as the instruments in the first stage cubic polynomials in age, education, fixed effects, and their interactions.

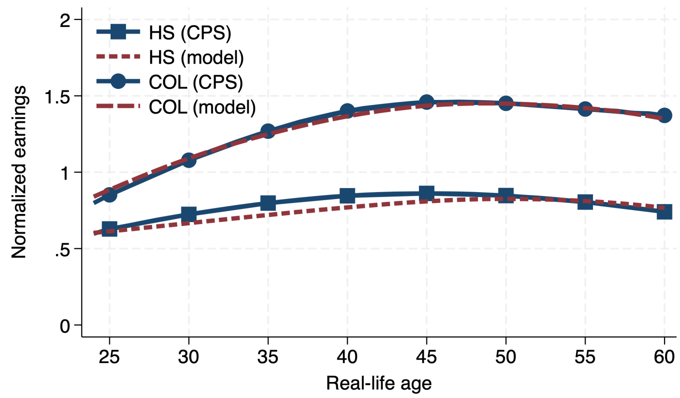

Figure 2: Pre-government earnings in the model and CPS, by age and education.
Notes: The profiles are normalized by the average economy-wide pre-government earnings.

empirical profiles in Figure 2, we regress log hourly wages on the household head's age, age squared, and education. It is clear from the figure that our human capital mechanism can replicate very closely the earnings profiles of both education groups.

Persistent shock. The idiosyncratic productivity shock y_j follows an AR(1) process given by (3). Under this specification, the conditional variance of y_j increases with age:

$$Var(\log(y_j)) = \sigma_\epsilon^2 \times \sum_{m=0}^{j-2} \rho_y^{2m}. \tag{28}$$

Following Storesletten et al. (2004), the parameters $(\rho_y, \sigma_\epsilon^2)$ are estimated outside the model by fitting (28) to the profiles of residual wage dispersion by age in the CPS. To construct the empirical variance profiles, we first compute the variance of the logarithm of the households' hourly wages by age and cohort and then remove cohort effects by regressing the computed variances on age and cohort dummies. Figure 3 plots the raw and the fitted coefficients on the age dummies, where the coefficients are normalized such that the variance is equal to 0 at age 25. This simple statistical model of earnings risk fits remarkably well the large increase in cross-sectional dispersion of log wages by age.

Fixed effect. We assume that the fixed effect v can take two values (v_L, v_H). We calibrate these values inside the model to match the Gini index of 0.40 for pre-tax earnings.

Production. The capital share ϖ and the depreciation rate δ are set outside the model to match the average ratio of capital income to GDP and the average ratio of investment to GDP. In line with Katz and Murphy (1992), ρ in (16) is set to 0.285, implying an elasticity of substitution between college and high-school graduates of 1.4. The TFP factor Z is normalized such that the weighted average of w_H and w_L is 1.0.

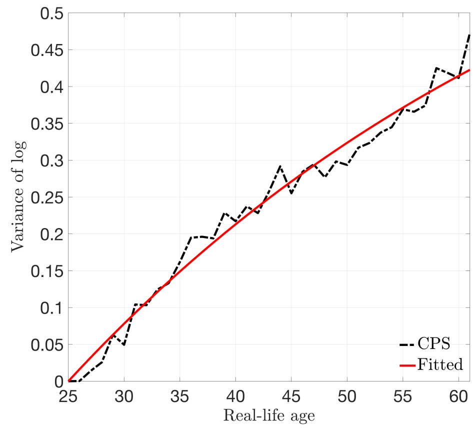

Figure 3: Cross-sectional variance of log hourly wages (CPS).

Notes: The variance is computed net of cohort effects via a cohort and age dummy-variable regression as in Deaton and Paxson (1994). The dashed line shows the coefficients on the estimated age dummies. The solid line represents the fitted variance according to the stochastic process in (3).

Social Security. The pension progressivity λ_{SS} is estimated outside the model by fitting the schedule of average replacement rates in (5) to its empirical counterpart. We calibrate the replacement rate level $\bar{\lambda}_{SS}$ inside the model to match the social security tax rate of 10.6% in the data. The bend point $\bar{e}_{\min}$ is set to 11% of the economy-wide average taxable earnings, consistent with the data. The parameter cap is calibrated inside the model by matching the fraction of households in the CPS whose pre-tax earnings exceed the maximum taxable earnings threshold (8%). Figure 4 shows that the specified schedule of replacement rates fits the data accurately. The penalty factor for early retirement δ_p in (6) is estimated inside the model by targeting the share of households at age 62 in the CPS who report receiving pension benefits (26%). Appendix B.3 provides more details.

Income tax-transfer program. We set the income tax progressivity λ_I to 0.216, which is the estimate computed by Heathcote et al. (2020) based on the Congressional Budget Office data. The equilibrium income tax level $\bar{\lambda}_{I,t}$ is equal to 0.118, which implies that a household whose taxable income is the same as the economy-wide average taxable income faces an average income tax of 11.8%. The share of wasted spending to GDP, $g_y = G/Y$, is equal to 7.8%, and the share of federal government debt to GDP, $d_y = D/Y$, is equal to 100%.[17] Finally, the consumption tax, τ_c, is set to 4.1% following Wu (2021).

[17] The dataset is provided by the Office of Management and Budget at the White House and is publicly accessible at `https://www.whitehouse.gov/omb/historical-tables/` (Table 7.1). Our model directly accounts for transfers through the income tax-transfer system and social security through the public pension system. Hence, to compute the empirical counterpart of G in the model, we sum up federal mandatory and discretionary government outlays, net of 1) transfers, such as outlays on income security, family assistance, and EITC, and 2) outlays on social security.

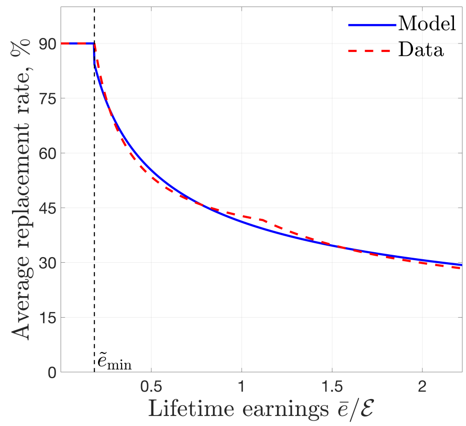

Figure 4: Estimated statutory schedule of average replacement rates.

Notes: The dashed vertical line marks the first bracket of the statutory replacement rate schedule. The average lifetime earnings $\bar{e}$ in the model and data are normalized by the economy-wide average taxable earnings $\mathcal{E}$.

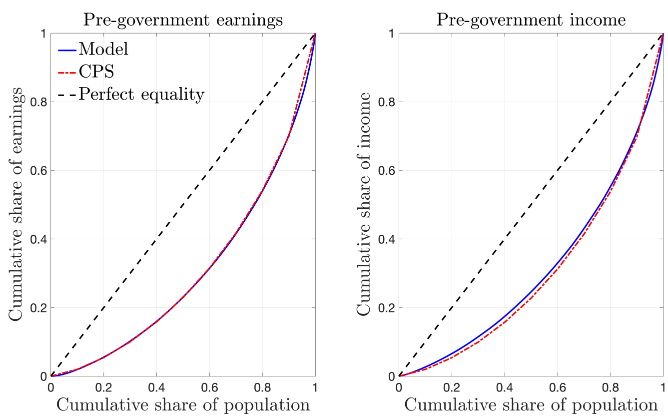

Figure 5: Earnings and income inequality in the model and data (CPS).

Notes: In the model, the Lorenz curves are constructed using the cross-sectional distribution of earnings and incomes in the initial steady state. The samples include households aged 25–62.

4.1 Untargeted moments

This section evaluates the model performance by looking at various moments that were not targeted in the calibration. Since one of the main channels through which the tax and social security systems shape social welfare is through the redistribution of earnings and wealth within and between generations, it is key that we have a reasonably good fit for the distribution and age profiles of wealth and earnings.

We first look at measures of earnings, income, and wealth inequality. Figure 5 shows that the model accurately fits the inequality in pre-government earnings and incomes as measured by the Lorenz curves. The model counterparts are constructed using the model-generated cross-sectional distributions of earnings and income in the initial steady state.

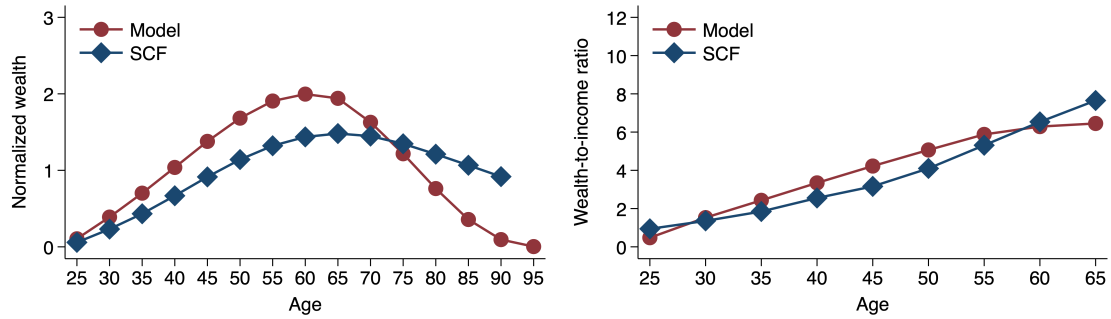

(a) Age profile of wealth (b) Age profile of wealth-to-income ratios

Figure 6: Untargeted moments of the wealth distribution.

Notes: The empirical profiles in both panels are constructed using SCF 2010–2019. The values on the horizontal axis represent the midpoint of 5-year age groups.

Figure 6 compares the fit of the age profiles of wealth (Panel a) and wealth-to-income ratios (Panel b). To construct the empirical profiles in both panels, we use data from the Survey of Consumer Finances (SCF) spanning the years 2010 to 2019. For the wealth profile in Panel (a), we follow the SCF definition of net worth, which is the total value of all assets net of all debts, and set negative wealth values to zero. The household wealth is normalized by the total wealth both in the model and the data. For the wealth-to-income ratios in Panel (b), we define income as the sum of wage income, income from sole proprietorships, farms, or other businesses and investments, and interest income. We construct the ratios at the household level, winsorize the top 2% of the ratios to remove outliers, and then compute the average ratio by age.

As the figure reflects, both the age profiles of wealth and the wealth-to-income ratios are matched quite well at most ages, even though they were not targeted. Note that our model is missing some potentially relevant features that could explain some of the discrepancies between the model and the data. One of these features could be home ownership. Real estate is a large part of wealth owned by households, and it is widely understood that home ownership both provides additional utility and is costly to adjust. Both of these factors could contribute to slower depletion of wealth during retirement.[18]

Next, we look at how the model performs in terms of lifetime earnings inequality. Using the U.S. Social Security Administration (SSA) records over the last six decades, Guvenen et al. (2022) document several features of real average lifetime earnings and lifetime earnings inequality since the 1970s. Two common indicators of inequality are the P90/P50

[18]In a follow-up paper (Ábrahám et al. 2025), we use a variant of this model without the detailed modeling of public pensions, to show that including housing and private retirement accounts can improve the fit of post-retirement wealth dynamics.

and the P50/P10 ratios of the (cross-sectional) lifetime earnings distribution, measuring inequality above and below the median, respectively. Assuming that agents entering retirement in our model are roughly the cohort aged 25 in 1977 in the data, our model does quite a nice job of matching these ratios. More specifically, the P90/P50 ratio of the distribution of lifetime earnings reported by Guvenen et al. (2022) among that cohort is around 2.45 for all cohorts (including men and women), while it is 2.0 in our model. As for the P50/P10 ratio, it is around 2.9 both in the data and the model.

We also explore the model fit regarding mortality rates by income and age, as well as the distribution of bequests. Appendix Figure 17 shows the model fit for mortality rates by income quintiles at different ages. While this was targeted in the calibration for age 70 (last panel), the remaining panels of the figure show that the model also fits the mortality rates by income at ages 40, 50, and 60, even though they were not targeted. Appendix Figure 19 shows the model fit for the distribution of bequests. The calibration has only targeted the share of deceased agents who leave no bequests, which is equal to 20%, but the figure shows a good match of the overall distribution as well.

5 Quantitative Results

When designing an optimal policy, the policymaker will face the classic tradeoff of public economics: efficiency versus redistribution/insurance. In our framework, redistribution is valued along two dimensions: within and between generations. Moreover, the policy instruments we consider distort labor supply (and human capital accumulation), as well as asset accumulation. The policymaker will therefore try to limit these distortions while simultaneously delivering desirable levels of redistribution, a tradeoff that will be the main theme of our quantitative analysis. Throughout the section, we assume that the policymaker maximizes the utilitarian welfare of currently alive generations subject to no constraints limiting redistribution. In Section 6, we conduct robustness exercises with respect to these two assumptions by (i) imposing constraints that limit redistribution and (ii) maximizing an alternative objective function.

To be able to relate more closely to the literature on optimal tax progressivity and optimal pension progressivity, the first subsection starts by studying cases in which the only policy instruments are pension and income tax progressivity, either together or in isolation. This will allow us to understand the distinct redistributive and distortionary properties of these two policy instruments, which turn out to be the key to understanding the more complex reforms discussed in the second subsection, where we allow the policymaker to optimize also over the generosity of pensions.

5.1 Optimal tax and pension progressivity

The results of the experiments in which the policymaker can optimize over income tax and/or pension progressivity are displayed in Table 2. The table reports the optimal policies, the average welfare gains (in consumption equivalent terms, in %) of the generations alive at the time of the reform (*Alive*) and the newborns (labor market entrants) in the final steady state (*Future*), as well as the political support (in %) among the currently alive.[19] The row labeled *Status quo* refers to the calibrated U.S. policy parameters.

Table 2: Optimal progressivity policy and welfare results.

	Policy			CEV, %		Support,
	$\bar{\lambda}_{SS}$	λ_{SS}	λ_I	Alive	Future	%
Status quo	0.413	1.420	0.216	–	–	–
Optimal income tax progressivity	0.413	1.420	0.165	0.208	0.256	59.1
Optimal pension progressivity	0.413	0.400	0.216	0.672	0.121	63.6
Alternative pension progressivity reform	0.413	2.100	0.216	0.641	-1.579	37.9
Optimal joint progressivity	0.413	2.100	0.140	1.399	0.431	61.8

Notes: The policy parameters include pension progressivity λ_{SS} and/or income tax progressivity λ_I. *Status quo* is the calibrated policy in the pre-reform steady state. The welfare gains are expressed in consumption equivalent terms (CEV, %) and reported for currently alive agents (*Alive*) and newborns in the final steady state (*Future*). The last column shows political support among alive agents.

The second row after *Status quo* displays the results for the case in which the government can only optimize over income tax progressivity, keeping the two pension policy parameters fixed. Compared to the status quo, the income tax reform prescribes a considerable decrease in progressivity (λ_I moves from 0.216 to 0.165), which generates positive but modest welfare gains for both the current and future generations of 0.21% and 0.26% in consumption equivalent terms, respectively. To better understand the source of these welfare gains, the solid line in Figure 7 plots the transitional dynamics of some of the key macroeconomic aggregates and tax rates from the pre-reform steady state to the final steady state associated with the optimal policy. As shown by the figure, a key source of welfare improvements are the lower aggregate distortions. A lower tax progressivity leads to a significant increase in effective labor, investment, and output at impact, as well as a sizable increase in capital accumulation, labor supply, and output in the long run.

Another important observation is the fact that the lower distortions come at the ex-

[19]Currently alive generations comprise all agents aged 25–100 at the time of the reform. Their total welfare, given by (22), takes into account the transitional dynamics from the pre-reform to the final steady state associated with the optimal solution.

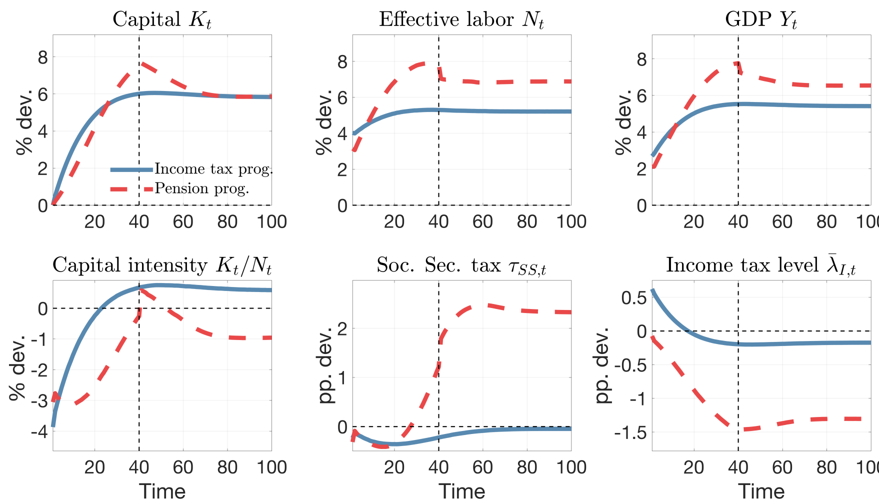

Figure 7: Transitional dynamics for the optimal progressivity reforms.

Notes: The reform maximizes the welfare of all currently alive generations. The paths are depicted for two scenarios: 1) income tax progressivity changes to its optimal level, with all other instruments remaining unchanged, and 2) pension progressivity changes, with all the other instruments remaining unchanged. The dashed vertical lines mark the completion of the phase-in period for the pension reform.

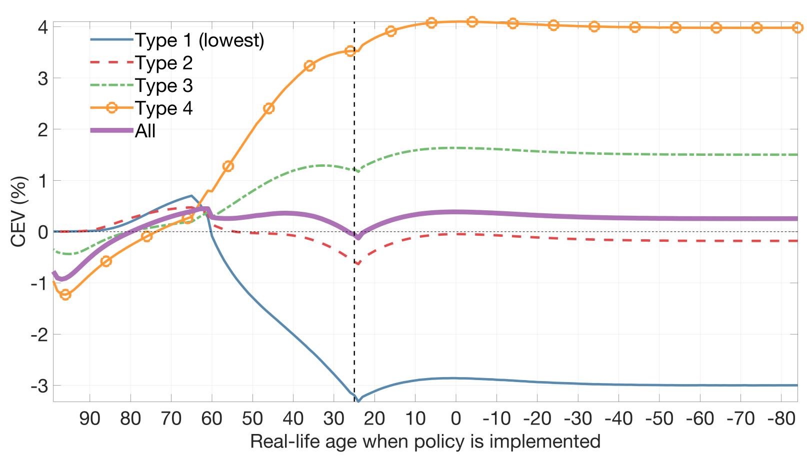

Figure 8: Welfare effects by type and age in the tax progressivity reform (CEV, %).

Notes: The reform maximizes the welfare of all currently alive generations. The vertical dashed line denotes labor-market entrants at the time of reform implementation (real-life age 25). Types are defined by education level z and fixed effect v as follows: Type 1 ($z = z_L, v = v_L$), Type 2 ($z = z_H, v = v_L$), Type 3 ($z = z_L, v = v_H$), and Type 4 ($z = z_H, v = v_H$).

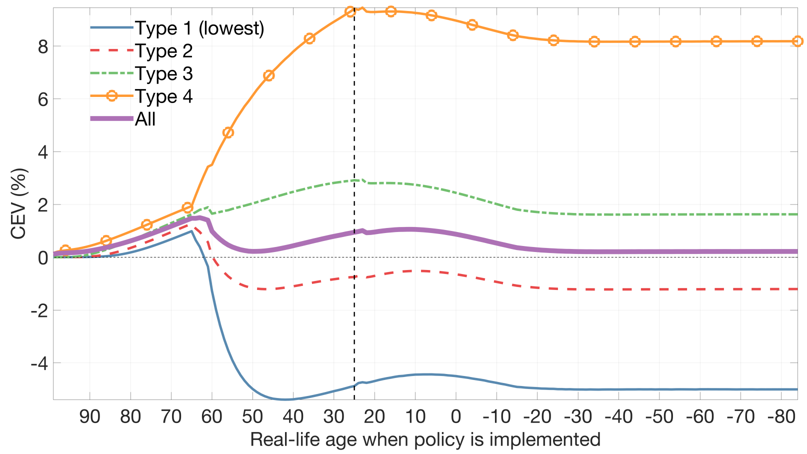

Figure 9: Welfare effects by type and age in the pension progressivity reform (CEV, %).
Notes: See notes to Figure 8.

pense of increased inequality among the current and future working households. This is illustrated in Figure 8, which depicts the welfare gains for the currently alive and future generations depending on their permanent income type, together with the average welfare gains. In the figure, the currently alive are located to the left of the vertical line, which corresponds to labor-market entrants at the time of the reform (real-life age of 25), while cohorts to the right of the vertical line are the children and the unborn generations.[20] While there are positive welfare gains (on average) for both the current and future generations, the policy change benefits higher-income types (Type 3 and 4), and it hurts agents with relatively low incomes (Type 1 and 2). These results are somewhat surprising, as our model features many key aspects of the high inequality in the U.S., and the utilitarian welfare function (ceteris paribus) favors redistribution towards low-income (low consumption) households. At the same time, our framework exhibits high levels of distortions, and the welfare benefits of reducing these outweigh the welfare costs of a higher inequality.[21]

Next, we explore whether optimizing only over pension progressivity can deliver a better outcome. The results for this case are depicted in the third row of Table 2, illustrating that the qualitative features of the optimal policy are similar to those in the previous reform. First, the optimal policy prescribes a substantial decrease in pension progressivity (λ_{SS} moves from 1.42 to 0.4), leading again to welfare gains for both current and future

[20]To calculate the welfare effect for newborns along the transition and in the final steady state, we use the distribution from the initial steady state, $F_{0,1}$. Note that $F_{t,1}$ varies over time only because bequests are endogenous. Since the probability of receiving bequests is tiny for newborns (Appendix Figure 18), the assumption of using weights from the initial steady state makes no material difference.

[21]These results are consistent with Heathcote et al. (2017), who find that the efficiency benefits of a lower tax progressivity outweigh the redistributional costs in a setting with endogenous labor supply and human capital accumulation.

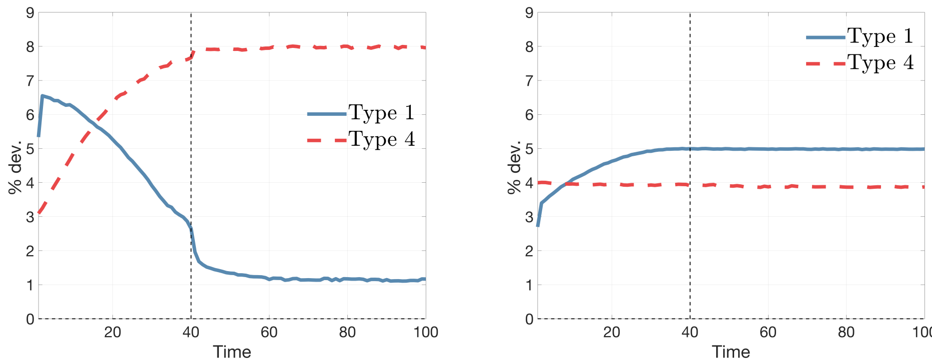

(a) Pension progressivity reform. (b) Income tax reform.

Figure 10: Transitional dynamics of hours worked by type in the pension progressivity and income tax progressivity reforms.
Notes: See notes to Figure 8.

generations. However, the gains for current generations are approximately three times higher than those in the previous experiment. Second, Figure 9 shows that the policy change also leads to an increase in inequality, which is also larger than the one generated in the income tax progressivity case, with the richest type gaining more than 8% and the poorest type losing more than 4% in the long run, as opposed to gains of 4% for the high-income types and losses of 3% for the low-income types in the income tax reform.

To understand why inequality increases, note that under the new optimal value of $\lambda_{SS} < 1$, the average replacement rate increases with the agent's average lifetime earnings, making the pension system regressive. This is illustrated in Figure 1, which displays the statutory schedules of the average replacement rates for different values of the pension policy parameters.[22] Given this, pensions increase substantially for high-income types, while they decrease significantly for low-income types, and it is therefore not surprising that inequality increases across all (working) currently living and future generations. In what follows, we explore why this shift in the pension system is capable of delivering significantly higher welfare gains despite the substantially higher inequality.

The dashed lines in Figure 7 depict the aggregate paths for the main macroeconomic variables and taxes for this second policy experiment. As in the previous case, the optimal policy reduces aggregate distortions. Importantly, however, the lower pension progressivity generates larger increases in both aggregate labor and output than the income tax

[22]To put the shift in pension progressivity λ_{SS} from 1.42 to 0.40 into perspective, note that—holding generosity $\bar{\lambda}_{SS}$ fixed—it implies double the replacement rate for an earnings-rich worker ($\bar{e}/\mathcal{E} = 2$) and half of the replacement rate for an earnings-poor worker ($\bar{e}/\mathcal{E} = 0.5$).

progressivity reform (except for the very first periods). One of the key reasons is the fact that the declining replacement rate schedule in the status quo acts as an implicit progressive tax on labor income, since an extra hour of work raises future pension benefits less as income increases. In contrast, with a regressive pension system, the implicit tax becomes effectively regressive.

It turns out that this shift in implicit taxation has very important consequences for the allocation of labor supply across income types. To see this, Figure 10 depicts the percentage deviation in labor supply relative to the status quo for the highest (Type 4) and lowest (Type 1) permanent income types. First, the long-run increase in labor supply is mainly coming from high-productivity households, in line with the change in implicit taxes. In addition, in the short run, low-income agents also tend to increase labor supply considerably. As long as taxes remain low, these agents find it optimal to work more and offset lower future pensions by increasing lifecycle savings. Given all this, the reduction in pension progressivity considerably improves the allocation and composition of labor supply, boosting labor supply both in the short and the long run. Importantly, this reshuffling of labor supply is not present in the income tax progressivity case, as shown in the right panel of the same figure. In this case, the labor supply of the two types changes very similarly, and if anything, the low-productivity type increases labor supply slightly more in the long run, a key reason why reducing pension progressivity leads to a significantly higher rise in effective labor and output in the second reform.

Another important difference between the two reforms is the fact that lowering pension progressivity triggers significant changes in both social security taxes τ_{SS} and average income taxes $\bar{\lambda}_I$, while the income tax progressivity reform has a much smaller impact (Figure 7). On the one hand, a regressive pension system with a higher replacement rate for the rich is more costly to run, and the social security tax rate needs to increase. On the other hand, the income tax base is expanding, both because households (especially high-productivity households) work more, and because the pensioners who used to be high-productivity workers receive more generous pensions. As a result, the average income tax rate $\bar{\lambda}_I$ can drop substantially. For low-income households, these two tax changes partially offset each other. However, for a significant proportion of high-type households, only the reduction of average income taxes matters, since their earnings exceed the cap beyond which they have to pay social security taxes. In other words, their "marginal effective" tax rate is not affected by the increases in τ_{SS}. This also explains why the high-productivity agents gain significantly more under this reform, while the welfare losses of low-income households, whose pensions drop substantially, are not that different.

Given the previous results, one could conclude that utilitarian social welfare in our

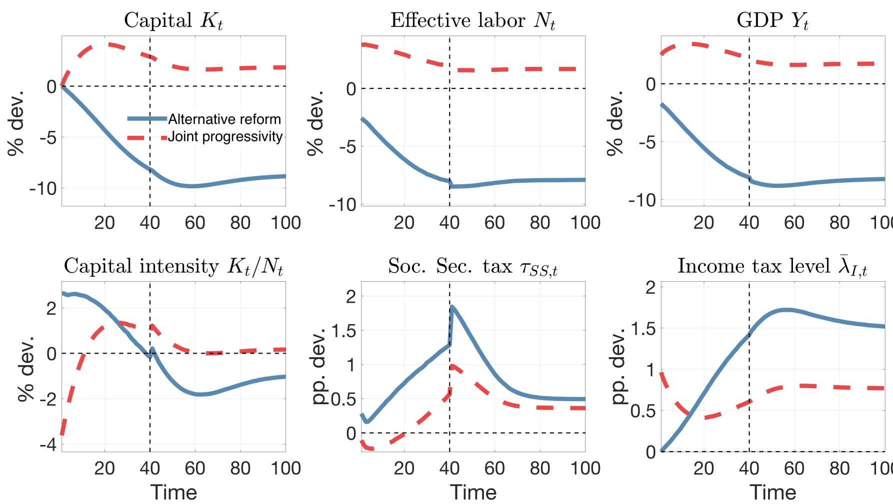

Figure 11: Transitional dynamics for alternative pension progressivity and joint optimal progressivity reforms.

Notes: The reform maximizes the welfare of all currently alive generations. The paths are depicted for two scenarios: 1) alternative pension progressivity reform and 2) joint income tax and pension progressivity reforms. The dashed vertical lines mark the completion of the phase-in period for the pension reform.

economy can be improved by decreasing the progressivity of either the tax or the pension system to reduce aggregate distortions. Interestingly, however, a last important difference between the previous two experiments is that, when the policy choice is pension progressivity, there is a second (local) maximum that yields very similar welfare gains for the currently alive but drastically different equilibrium and inequality properties.

The results for this case are displayed under "Alternative pension progressivity reform" in Table 2, while the aggregate paths are displayed in Figure 11 (solid line) and the welfare effects are shown in Figure 12. As opposed to the previous reforms, this case prescribes a significantly higher pension progressivity compared to the status quo (λ_{SS} moves from 1.42 to 2.1), leading to a higher implicit tax progressivity and to a considerable decrease in inequality, with the lowest income type (Type 1) experiencing large welfare gains, while the highest type (Type 4) experiences large welfare losses. At the same time, the solid line paths in Figure 11 illustrate that the decrease in inequality comes at the expense of an increase in distortions and a smaller economy both in the short and the long run, resulting in a sizable welfare loss for future generations (1.58 %), in contrast to the previous progressivity reforms, which lead to welfare gains for both future and current generations.

Note that, for current generations, the distortions are not as costly due to the phasing-in of the pension reform. First, social security taxes need to increase only gradually during the phase-in period, while they decrease considerably after, as lower pensions will need

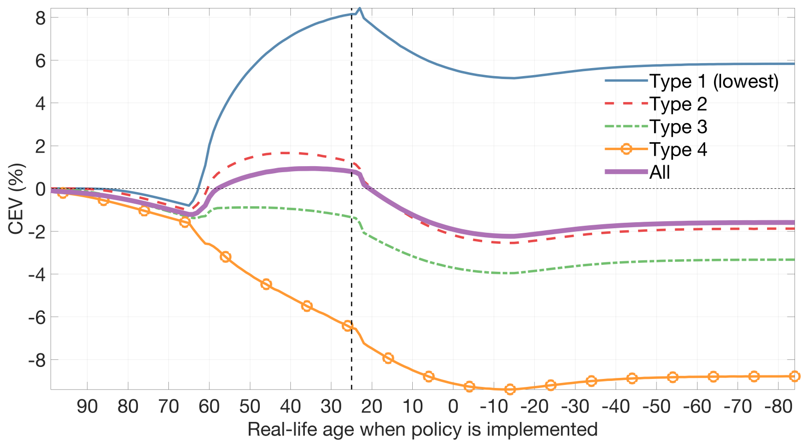

Figure 12: Welfare effects by type in alternative pension progressivity reform (CEV, %).
Notes: See notes of Figure 8.

to be paid out to higher-income agents. In addition, lower capital accumulation and labor supply lead to a lower tax base, so that income taxes have to remain high in the long run, but they increase only gradually. Last, the losses of high-productivity agents are more than compensated by the high gains of the low-income ones.[23]

Importantly, such a local maximum does not exist when the only policy instrument is income tax progressivity. In this case, the only way to improve on the status quo is by reducing income tax progressivity and hence by increasing inequality. Under the status quo policy, the marginal cost of distortions when increasing tax progressivity is higher than the marginal benefit of reducing inequality. Also, as we reduce income tax progressivity, the marginal social benefits from reduced distortions are getting smaller, while the marginal social costs of increased inequality are increasing, generating a unique global optimum. While a similar intuition holds when we consider a reduction in pension progressivity with respect to the status quo, this is not the case with an increase in pension progressivity. In the short and medium run, which is key for the welfare of the currently living generations, the distortions do not increase that much, as the new pension system is only applied to an (increasing) share of households, leading to a local optimum.

So far, our isolated reforms illustrate that there is a key conflict between (output) efficiency and (within-cohort) inequality. When we reduce distortions, we need to increase inequality, and vice versa. Next, we show that allowing the social planner to jointly choose

[23]Note also that pensioners suffer welfare losses when income taxes are increasing and/or interest rates drop during their remaining life span. The tax increases in Figure 7 are responsible for the welfare loss of pensioners in the optimal tax progressivity reform, while the drop in the interest rate (due to the increase in K/N), together with the gradual tax increases depicted in Figure 11 cause pensioners to suffer welfare losses in the reform that increases pension progressivity.

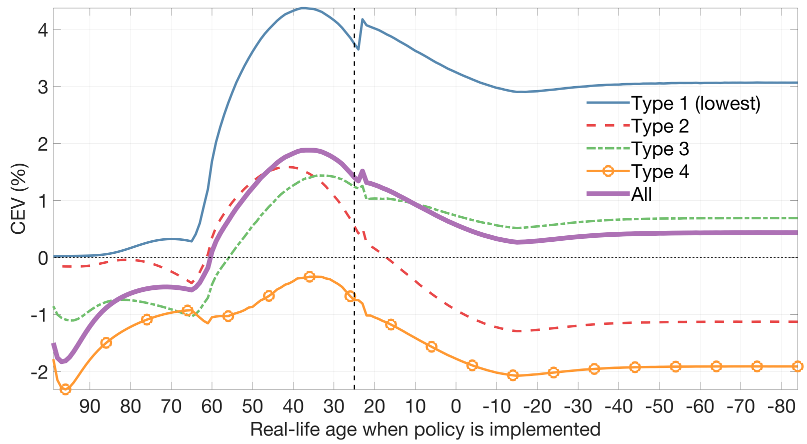

Figure 13: Welfare effects by type and age in the joint progressivity reform (CEV, %).
Notes: See notes of Figure 8.

pension progressivity and income tax progressivity can largely eliminate this conflict. The results of this reform are depicted in the last row of Table 2, while the aggregate paths and welfare gains are depicted in Figures 11 (dashed line) and 13, respectively. As evident from the table and the figures, the joint optimal policy is a clear winner, with much higher welfare gains for the currently alive generations (1.4%) than in the previous cases and non-trivial welfare gains for the future ones (0.43%). In terms of the redistribution outcomes, the increased pension progressivity dominates, and the policy change generates a smaller inequality. In terms of distortions, the reduction in tax progressivity dominates, and there are higher savings, labor supply, and output both in the short and long run.

The key to higher welfare gains is that the welfare losses of the high-income individuals are contained through the lower tax progressivity. In turn, the better provision of labor supply and saving incentives in the combined reform can be seen in the dashed line of Figure 11. Comparing the joint and alternative pension reforms isolates the impact of reducing tax progressivity on all variables, since this is the only policy that is different across the two reforms. Lowering income tax progressivity implies lower tax changes and altogether smaller aggregate distortions. This also solves the intergenerational conflict, as all cohorts (except the pensioners) gain, as shown by the line "All" in Figure 13.

To sum up, we have shown that adjusting pension progressivity may bring much larger welfare gains than optimizing over income tax progressivity. Furthermore, as opposed to tax progressivity, we can obtain significant welfare gains by either reducing or increasing pension progressivity, with very different redistributional consequences both within and

Table 3: Optimal full reform and welfare results.

	Policy			CEV, %		Support,
	$\bar{\lambda}_{SS}$	λ_{SS}	λ_I	Alive	Future	%
Status quo	0.413	1.420	0.216	–	–	–
Optimal full reform	0.800	0.400	0.191	7.929	-5.673	89.3
Alternative pension reform	0.717	2.100	0.165	6.172	-4.709	84.3

Notes: The policy parameters include pension progressivity λ_{SS}, income tax progressivity λ_I and pension generosity $\bar{\lambda}_{SS}$. "Alternative pension reform" imposes the constraint that no Type 1 agents—whether currently alive, along the transition path, or in the final steady state—are made worse off (on average). See also notes in Table 2.

across generations.[24]

In the next section, we show that these findings remain largely robust when we include the generosity of the pension system as an additional policy instrument.

5.2 Optimal full reform with pension generosity

So far, we have kept pension generosity $\bar{\lambda}_{SS}$ fixed and this subsection extends the set of policy instruments to this instrument as well. The second row of Table 3 contains the optimal policy and welfare gains of the full optimal policy with pension generosity as an additional instrument.

There are three key observations. First, the optimal policy inherits the properties of the optimal pension progressivity reform, with a considerable reduction in pension progressivity that makes the system regressive (λ_{SS} moves from 1.42 to 0.40), while there is only a very modest change in income tax progressivity. Second, pension generosity $\bar{\lambda}_{SS}$ increases substantially in both cases. Of course, this considerably improves the welfare of current generations compared to the previous reforms (7.9% in the optimal policy and 6.2% in the locally optimal policy). Third, a higher generosity increases distortions significantly, generating large welfare losses for future generations of around 6.7%.

Next, we elaborate more on the role of pension generosity. An important observation is that increasing pension generosity reduces lifecycle savings incentives and, hence, leads to an economy with lower aggregate capital and output in the long run. However, this is a gradual process due to the phase-in, implying that the current generations can still benefit significantly.

[24] As we have explained in the introduction, the relatively small literature on optimal pension progressivity is inconclusive on whether progressivity should increase or decrease compared to the status quo. Our paper shows that one can improve welfare in both ways. In Section 6.3, we will also show that either policy can be globally optimal depending on the severity of distortions.

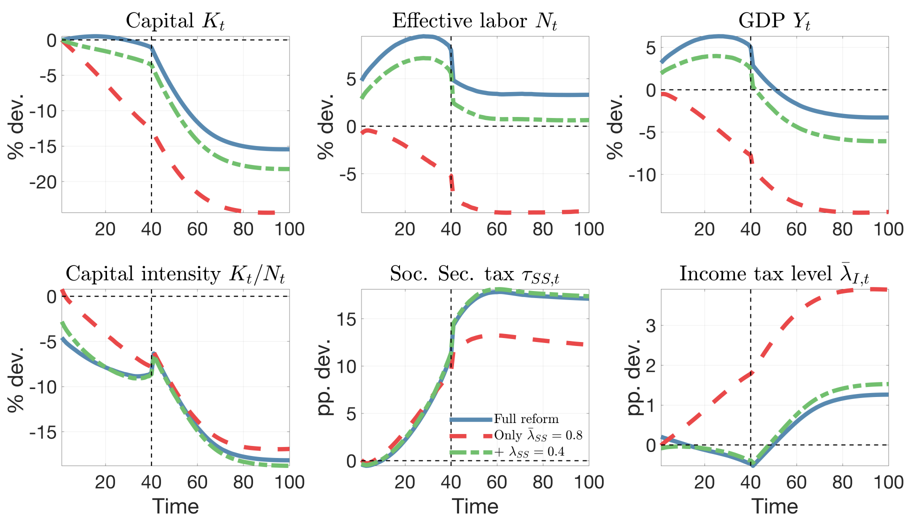

Figure 14: Decomposition of transitional dynamics for the full optimal reform.

Notes: The unconstrained reform maximizes the welfare of all currently alive generations. The paths are depicted for three scenarios: 1) pension generosity changes to its optimal level, with all other instruments remain unchanged (*Only* $\bar{\lambda}_{SS}$), 2) pension generosity and progressivity change, and income tax progressivity remains unchanged ($+\lambda_{SS}$), and 3) all policy instruments change (*Full reform*). The dashed vertical lines mark the completion of the phase-in period for the pension reform.

Figure 14 plots the transitional dynamics when we implement the reform in three steps: first, we only set the pension generosity $\bar{\lambda}_{SS}$ to its optimal level, we then add the optimal pension progressivity λ_{SS} and, finally, we incorporate the optimal income tax progressivity λ_I, obtaining the full reform. In line with the intuition above, if the policymaker only increases pension generosity, both aggregate capital and labor supply drop dramatically in the long run. At the same time, the reduction in pension progressivity (and tax progressivity) offsets this drop in terms of labor supply, and it mitigates the decrease in long-run capital accumulation. This highlights again the crucial role of pension progressivity in reducing aggregate distortions.

Notice that in the short term, the capital stock remains almost unchanged, both labor and output rise, and social security taxes increase only gradually at the beginning, while they rise dramatically in the long run to finance higher pensions. This differential impact of the reform in the short run versus the long run is crucial for understanding why currently alive generations gain so much, while future generations lose (see the welfare effects by age and type in Appendix Figure 21).

The third row of Table 3 displays the results of an alternative pension reform that is designed to reduce inequality, which we compute by imposing the constraint that no Type 1 agents of any generation are worse off after the policy change compared to the status quo. As one can see, this policy still delivers very high welfare gains to the current gen-

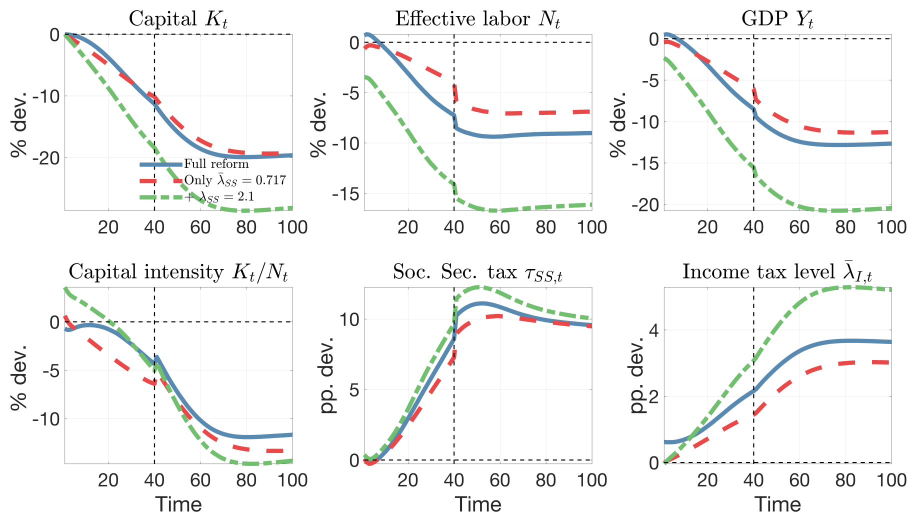

Figure 15: Decomposition of transitional dynamics for the *"Alternative pension reform"*.
Notes: The reform maximizes the welfare of all currently alive generations subject to the constraint that no Type 1 agents—whether currently alive, along the transition path, or in the final steady state—are made worse off (on average). Also see notes to Figure 14.

erations (6.172% compared to 7.929% in the optimal full reform). Hence, our previous finding showing that there are two different ways of delivering high welfare gains when optimizing over pension progressivity extends when the policymaker can also use pension generosity as a policy instrument.

In the alternative pension reform, pension generosity still increases, but the prescription for pension progressivity is drastically different from the other reform, with a large increase in progressivity that has the reverse impact on within-cohort inequality (Appendix Figure 22). In Figure 15, we decompose the transitional dynamics for this reform. As before, the redistribution gains of a higher pension progressivity come at the cost of reduced efficiency, with a much higher decrease in output (20% versus 4% in the full optimal reform) and a much larger decrease in capital and effective labor due to the increase in both pension progressivity and generosity. Current generations still benefit from this policy, as the negative consequences of higher taxes ($\tau_{SS,t}$ and $\lambda_{I,t}$) and lower wages are delayed, and because many of them will enjoy higher pensions. It is also clear from the decomposition that the reduction in income tax progressivity that accompanies the increase in pension progressivity is critical to achieve these welfare gains, as it mitigates the distortions considerably, reducing the drop in labor by almost a half. Overall, this shows that our main findings still persist when adding generosity as a policy instrument.

5.3 Main lessons from the quantitative analysis

To summarize, our quantitative analysis yields the following key findings. First, pension progressivity is a powerful instrument for both managing distortions and redistributing within generations, even when adjusting income tax progressivity or pension generosity are also instruments available for the social planner. As the existing literature has mostly concentrated on adjusting the tax-transfer system to determine the optimal balance between redistribution and output efficiency, this is a new and important insight.

Second, there are potentially two very distinct ways to improve social welfare when the social planner can jointly redesign the tax-transfer and the pension systems. The first focuses on reducing distortions and features regressive pensions, resulting in higher inequality. The second focuses on reducing inequality and hence features very progressive pensions complemented with a much less progressive income tax system. Both of these reforms would gain overwhelming political support among the currently living generations, although the identity of the main winners would be drastically different. For this reason, the choice between these policies is a political question.

In the next section, we conduct a series of sensitivity checks to figure out whether our key results remain robust.

6 Robustness

In this section, we check the robustness of our key results in the full reform of Section 5.2 to three different alternatives. First, we consider constraints that limit redistribution between generations, as both the globally and locally optimal reforms in Table 3 generate vast redistribution from future to current generations. After that, we address the same issue by using a social welfare function that puts all the weight on future generations. Finally, we analyze (counterfactual) versions of our environment in which some of the channels creating distortions are eliminated.

6.1 Limiting redistribution between generations

In this subsection, we study the properties of the optimal policy when there are limits on intergenerational redistribution. Specifically, we look for the policy that maximizes the welfare of currently alive generations subject to the constraints (24) and (25), which ensure that neither currently alive cohorts nor any of the generations along the transition or in the final steady state are made worse off on average after the reform. The resulting reform is labeled "*No cohort is worse off*" in Table 4.

Table 4: Optimal constrained policy and welfare results.

	Policy			CEV, %		Support,
	$\bar{\lambda}_{SS}$	λ_{SS}	λ_I	Alive	Future	%
Status quo	0.413	1.420	0.216	–	–	–
No cohort is worse off	0.488	0.753	0.214	1.913	0.0	77.2

Notes: The policy parameters include pension generosity $\bar{\lambda}_{SS}$, pension progressivity λ_{SS}, and income tax progressivity λ_I. "No cohort is worse off" imposes the constraint that no currently alive cohorts and any generation along the transition and in the final steady state is made worse off on average. See also notes to Table 2.

The qualitative features of the optimal pension system remain the same as in the full optimal reform: pension generosity increases and the pension system becomes regressive compared to the status quo. However, the magnitude of the changes in the policy variables is significantly smaller. First, we see a much smaller rise in pension generosity, with $\bar{\lambda}_{SS}$ moving from 0.41 to 0.49, while it almost doubled in the unconstrained optimal case. This is not surprising, since constraints (24) and (25) limit intergenerational redistribution, making it infeasible to increase pension generosity significantly. Second, the decline in pension progressivity, while also more moderate than in the full optimal reform, is still substantial, with λ_{SS} dropping from 1.42 to 0.75.

These results should not be surprising, as we have already seen in Figure 9 that reducing pension progressivity alone can improve welfare in a way that generates practically no intergenerational conflict, while we have learned in Section 5.2 that increasing generosity improves the welfare of the current generations. The results in this table show that a combination of a smaller rise in generosity and a more modest drop in pension progressivity can keep future generations indifferent and increase significantly the welfare of current cohorts (Appendix Figure 24). The key is that the savings and labor distortions induced by increasing pension generosity are more than offset by the reduction in pension progressivity, leading to an increase in output both in the short run and the long run (Appendix Figure 23). The key difference compared to the unconstrained optimal reform is that, in that case, generosity increases so much that an even larger reduction in pension progressivity can only mitigate but not offset the higher distortions.

We have already shown that combining a higher pension progressivity and a lower income tax progressivity can improve welfare for all generations except for the pensioners (Table 2 and Figure 12). Hence, the reforms that reduce inequality do not create intergenerational conflicts if we do not take into account the retirees (and households very close to the retirement age). It turns out that we cannot do significantly better even if we allow for

Table 5: Optimal policy and welfare results with alternative objective.

	Policy			CEV, %		Support,
	$\bar{\lambda}_{SS}$	λ_{SS}	λ_I	Alive	Future	%
Status quo	0.413	1.420	0.216	–	–	–
Optimal full reform	0.285	0.400	0.250	-2.428	1.026	43.1
No cohort is worse off	0.412	0.698	0.228	0.340	0.688	62.9
Alternative pension reform	0.354	2.100	0.132	0.096	0.692	43.4

Notes: In all experiments, the government maximizes the welfare of newborn agents in the final steady state. "No cohort is worse off" imposes the constraint that no currently alive cohorts and any generation along the transition and in the final steady state is made worse off on average. "Alternative pension reform" imposes the constraint that no Type 1 agents—whether currently alive, along the transition path, or in the final steady state—are made worse off (on average). See notes in Table 2.

changing generosity. In particular, there is no reform that reduces inequality and weakly increases the welfare of all cohorts, including the retirees. The main reason is that households that have already retired do not benefit from the increase in pension progressivity, but tend to lose from the reduction in income tax progressivity, as their income falls below the average income and reduced progressivity implies a higher tax burden. In sum, with this caveat, the main qualitative findings from the full unconstrained reforms remain valid, even if the policymaker has constraints on intergenerational redistribution.

6.2 Maximizing the welfare of future generations

In this subsection, we check whether our key results remain robust when the welfare weights shift fully towards future generations. To this end, we compute reforms that maximize the welfare of newborns in the final steady state under the same constraints as in the benchmark. The results are summarized in Table 5. Below, we highlight the main takeaways from this analysis, while Appendix D.2 provides a more detailed discussion.

Table 5 reveals a stark difference in optimal pension generosity compared to the benchmark: when the government assigns full weight to future generations, all reforms prescribe a decrease in generosity. This should not be surprising, as generosity distorts both capital accumulation and labor supply. Reducing it, therefore, leads to a larger economy in the long run, benefiting future generations at a cost to current generations, who must work more and consume less to build a larger capital stock. As in the benchmark, there is substantial intergenerational redistribution, although now in the opposite direction.

Importantly, the remaining qualitative (and quantitative) results remain in line with those in the benchmark. First, both the unconstrained and the *No cohort is worse off* reforms feature regressive pensions, since a lower pension progressivity leads to a significant drop in distortions and, consequently, to a larger economy that benefits most future agents. Second, these reforms lead to redistribution towards the income-rich and hence increase inequality (in welfare). Last, the *Alternative pension reform* shows that there is scope to reduce inequality through an increase in pension progressivity, which is compensated by a reduction in tax progressivity, generating 67% of the welfare gains in the unconstrained reform. In sum, the key insights remain robust regardless of whether the social planner prioritizes current or future generations. There are two distinct reform types able to generate substantial welfare gains, and pension progressivity plays a central role in both.

6.3 Counterfactual experiments

As we have seen, there exist reforms delivering high welfare gains by either reducing inequality or by reducing distortions, but the highest welfare gains are achieved by doing the latter. This section conducts counterfactual experiments that shut down important channels in our model that may affect the equity-efficiency tradeoff. In all experiments, the government maximizes the welfare of currently alive generations subject to no additional constraints on redistribution.

In a first set of experiments, we consider counterfactual versions of our economy in which some of the channels creating distortions are eliminated. In particular, we assume (i) an exogenous age profile of productivity instead of endogenous human capital accumulation, (ii) exogenous retirement, (iii) separate taxation of labor and capital income with capital being taxed at an exogenous rate, and (iv) partial equilibrium (open economy with a fixed interest rate). In the first case, all dynamic amplification effects through the learning-by-doing mechanism are eliminated. In the second case, one important adjustment margin of labor supply, namely, retiring early, is eliminated. In the third case, we reduce tax distortions on older and richer agents by excluding asset income from the income tax base. Finally, the last reform eliminates the feedback mechanism of the interest rate, which reinforced the direct effect of distortions in the benchmark case.

The results of these reforms are displayed in Table 6 in Appendix D.3. The key observation is that optimal policy now calls for an increase in pension progressivity, in contrast to the optimal (unconstrained) policy in the baseline model. As we have discussed above, the common characteristic of these cases is that pension and tax progressivity distort labor supply less than in the benchmark. In other words, whenever distortions become quantitatively less relevant, as in these cases, the policy that reduces inequality and allows for more distortions becomes unconstrained optimal.

At the same time, we have found that even in these cases, there exists a drastically different policy (not an optimal one) that also substantially increases social welfare by *reducing* pension progressivity—and hence distortions—at the expense of *increasing* inequality. We have confirmed this by computing the welfare gains in each of the four cases under the optimal policy from the unconstrained benchmark. The results are displayed in Table 7 in Appendix D.3. As the table reflects, the benchmark optimal policy still achieves substantial welfare gains in these four counterfactual economies. Hence, our previous results regarding the existence of two possible policies that deliver high welfare gains and the key role of pension progressivity are robust to important model assumptions, but when distortions become quantitatively less relevant, the policy that reduces inequality and allows for more distortions becomes unconstrained optimal.

It is important to note that, in all four counterfactual experiments above, as well as in the benchmark alternative policy, the motive to provide redistribution and insurance dominates the objective of reducing distortions. However, only in the counterfactual with exogenous human capital, the increase in pension progressivity and generosity is accompanied by an *increase* in income tax progressivity. This results in additional redistribution toward the poor, but it also generates higher distortions. Nevertheless, the distortions remain considerably smaller than those the same policy would have induced in the benchmark with endogenous human capital. As a result, the redistribution benefits stemming from the increase in tax progressivity outweigh the efficiency losses.[25]

Next, we conduct a second set of experiments to check the robustness of other features that could affect the equity-efficiency tradeoff. In particular, we assume that pension benefits are not taxed, as taxing pension benefits progressively may mitigate the redistributive effects of reducing pension progressivity. We also eliminate the deductions of social security contributions from the income tax base, as this feature may "compensate" households for higher social security taxes. Finally, we consider common mortality rates across types. In our benchmark, low-income agents have a shorter life span and hence enjoy public pension benefits for a shorter period. In principle, this heterogeneity could be behind the result that the public pension system redistributes resources towards richer households who have longer life expectancies at retirement. The results of these reforms are displayed in the last three rows of Table 6 in the appendix. The table shows that the results are qualitatively and quantitatively similar to the ones in the baseline model, implying that these model features do not significantly affect the redistribution and aggregate efficiency ben-

[25]We do not recalibrate the model but impose the profile of average hourly wages from our benchmark with endogenous human capital, so our counterfactual is not comparable quantitatively to those studies that typically find that more redistribution (less inequality) is beneficial (e.g., Conesa et al. 2009).

efits/costs of adjusting pension progressivity (and the other instruments). Similar to the benchmark and the other counterfactual reforms above, there is still a welfare-enhancing policy that reduces inequality and increases distortions, but it delivers lower welfare gains. This is confirmed when looking at the last three rows of Table 7, achieving very high welfare gains in the three counterfactual economies. Overall, the main results of the benchmark model remain robust, but the level of distortions may affect the social ranking between the efficiency-enhancing and the inequality-reducing optimal policies.

7 Conclusions

We study the optimal policy mix between income tax progressivity, pension progressivity, and pension generosity in a rich quantitative model that takes into account many important elements of the U.S. social security and income tax-transfer systems and features many channels through which both systems can distort the economy. Our model, calibrated to the U.S. economy, generates realistic inequality in income, wealth, and consumption. Hence, it is well-suited to study the equity-efficiency tradeoff that the design of social security and tax-transfer systems faces.

One of the key conclusions of our analysis is that there are two types of drastically different policies that can deliver significant welfare gains. The most prominent property of the first one is a dramatic reduction in pension progressivity, which reduces distortions and increases inequality. The key property of the second type of reform is a considerable increase in pension progressivity accompanied by a substantial reduction in income tax progressivity. This policy results in increasing distortions, with a significant reduction in inequality. Both of these policies increase welfare significantly because our calibrated economy features both large distortions and inequality. Nevertheless, reducing distortions consistently delivers higher (utilitarian) social welfare across all cases. We find that this is primarily due to the endogenous human capital and endogenous retirement channels in our baseline model. An important new insight from these results is that pension progressivity is a powerful instrument for either reducing labor supply distortions (first type of reforms) or redistributing towards the income-poor (second type). Our paper is the first to analyze the role of this instrument in tackling the equity-efficiency tradeoff in a state-of-the-art quantitative lifecycle model.

Our environment can be a starting point for several future projects. In our current setting, households use their savings to supplement public pensions for old-age consumption, so not distinguishing between regular assets and private pension funds is unlikely to affect our main findings. However, since private retirement plans serve as an important source of wealth and old-age consumption and also offer fiscal advantages over regular savings, one extension is to analyze the implications of different tax advantages of these plans for efficiency and redistribution. This is the avenue we pursue in Ábrahám et al. (2025), where we use a simpler model for public pensions for tractability. Another extension is to incorporate a demographic transition and examine which policy combinations would best mitigate the intergenerational redistribution caused by population aging.

References

Ábrahám, A., P. Brendler, and E. Carceles-Poveda (2025). Taxation of private retirement wealth. *Working paper*.

Ábrahám, A. and E. Carceles-Poveda (2008). Tax reform with endogenous borrowing limits and incomplete asset markets. *Working paper*.

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

Badel, A., M. Huggett, and W. Luo (2020). Taxing top earners: a human capital perspective. *The Economic Journal 130*(629), 1200–1225.

Bell, F. C., A. H. Wade, and S. C. Goss (1992). Life tables for the United States Social Security area. *Actuarial Study 116*, 1900–2100.

Benabou, R. (2000). Unequal societies: Income distribution and the social contract. *American Economic Review*, 96–129.

Brendler, P. (2020). Why hasn't Social Security changed since 1977? *Review of Economic Dynamics 36*, 134–157.

Brendler, P. (2023). Rising earnings inequality and optimal income tax and Social Security policies. *Journal of Monetary Economics 134*, 35–52.

Carroll, D. R., A. Luduvice, and E. Young (2023). Optimal fiscal reform with many taxes. Working Papers 23-07, Federal Reserve Bank of Cleveland.

Chetty, R., M. Stepner, S. Abraham, S. Lin, B. Scuderi, N. Turner, A. Bergeron, and D. Cutler (2016). The association between income and life expectancy in the United States, 2001-2014. *Journal of the American Medical Association 315*(16), 1750–1766.

Conesa, J. C. and C. Garriga (2008). Optimal fiscal policy in the design of social security reforms. *International Economic Review 49*(1), 291–318.

Conesa, J. C., S. Kitao, and D. Krueger (2009). Taxing capital? not a bad idea after all! *American Economic Review 99*(1), 25–48.

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

De Nardi, M. and F. Yang (2016). Wealth inequality, family background, and estate taxation. *Journal of Monetary Economics 77*, 130–145.

Deaton, A. and C. Paxson (1994). Intertemporal choice and inequality. *Journal of Political Economy 102*(3), 437–467.

Fehr, H. and C. Habermann (2008). Risk sharing and efficiency implications of progressive pension arrangements. *The Scandinavian Journal of Economics 110*(2), 419–443.

Fehr, H., M. Kallweit, and F. Kindermann (2013). Should pensions be progressive? *European Economic Review 63*, 94–116.

Fuster, L., A. Imrohoroglu, and S. Imrohoroglu (2007). Elimination of Social Security in a dynastic framework. *The Review of Economic Studies 74*(1), 113–145.

Fuster, L., A. Imrohoroglu, and S. Imrohoroglu (2008). Altruism, incomplete markets, and tax reform. *Journal of Monetary Economics 55*(1), 65–90. Carnegie-Rochester Conference Series on Public Policy: Advances in Dynamic Public Finance April 20-21, 2007.

Guner, N., R. Kaygusuz, and G. Ventura (2023). Rethinking the welfare state. *Econometrica 91*(6), 2261–2294.

Guvenen, F., G. Kaplan, J. Song, and J. Weidner (2022). Lifetime earnings in the United States over six decades. *American Economic Journal: Applied Economics 14*(4), 446–79.

Heathcote, J., K. Storesletten, and G. L. Violante (2017). Optimal tax progressivity: An analytical framework. *The Quarterly Journal of Economics 132*(4), 1693–1754.

Heathcote, J., K. Storesletten, and G. L. Violante (2020). How should tax progressivity respond to rising income inequality? *Journal of the European Economic Association 18*(6).

Heathcote, J. and H. Tsujiyama (2021). Optimal income taxation: Mirrlees meets Ramsey. *Journal of Political Economy 129*(11), 3141–3184.

Holter, H. A., D. Krueger, and S. Stepanchuk (2019). How do tax progressivity and household heterogeneity affect Laffer curves? *Quantitative Economics 10*(4), 1317–1356.

Huggett, M. (1996). Wealth distribution in life-cycle economies. *Journal of Monetary Economics 38*(3), 469–494.

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

Hurd, M. D. and J. P. Smith (2001). Anticipated and actual bequests. In *Themes in the Economics of Aging*, NBER Chapters, pp. 357–392. National Bureau of Economic Research.

Katz, L. F. and K. M. Murphy (1992). Changes in relative wages, 1963-1987: Supply and demand factors. *The Quarterly Journal of Economics 107*(1), 35–78.

Keane, M. P. (2011). Labor supply and taxes: A survey. *Journal of Economic Literature 49*(4), 961–1075.

Kindermann, F. and V. Pueschel (2021). Progressive pensions as an incentive for labor force participation. Technical Report 3928658, Rochester, NY.

Krueger, D. and A. Ludwig (2016). On the optimal provision of social insurance: Progressive taxation versus education subsidies in general equilibrium. *Journal of Monetary Economics 77*, 72–98.

Macnamara, P., M. Pidkuyko, and R. Rossi (2022). Taxing consumption in unequal economies. Economics discussion paper series, The University of Manchester.

MaCurdy, T. E. (1981). An Empirical Model of Labor Supply in a Life-Cycle Setting. *Journal of Political Economy 89*(6), 1059–1085.

Makarski, K., J. Tyrowicz, and O. Komada (2022). Progressing towards efficiency: the role for labor tax progression in reforming social security. *IZA Discussion Paper Series 15100*.

Nam, L. (2024). Optimal progressive pension systems and heterogeneous labor market risk. *Working paper*.

Nishiyama, S. and K. Smetters (2005). Consumption taxes and economic efficiency with idiosyncratic wage shocks. *Journal of Political Economy 113*(5), 1088–1115.

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

Nishiyama, S. and K. Smetters (2008). The optimal design of social security benefits. *Working Paper*.

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

Tran, C. and N. Zakariyya (2025). Progressive pension and optimal tax progressivity. *Working paper*.

Wu, C. (2021). More unequal income but less progressive taxation. *Journal of Monetary Economics 117*(C), 949–968.

Supplemental online appendix

A Model

A.1 Bequest distribution

This section discusses how the probability of receiving bequests and the bequest distribution ($\Gamma^{z,v}_{j+1,t+1}$ and $\Phi^{z,v}_{t+1}$) are determined in equilibrium. When we model bequests, we do not model families explicitly but instead introduce a dynamic link between generations in a stylized way through the transition matrix $\pi(z^c, v^c \mid z^p, v^p)$, where (z^p, v^p) and (z^c, v^c) refer to the parents' and children's productivity type, respectively.[26] In particular, when a household of type (z^p, v^p) passes away with positive wealth, $\pi(z^c, v^c \mid z^p, v^p)$ is the probability that their bequest is received by a type (z^c, v^c) household.

In the data, if a household has not received a bequest, we cannot distinguish whether this is because no one has died who could leave a bequest for that household or because the parents have passed away but have left zero bequests. Given this, we only keep track of bequests that are strictly positive. First, we show how $\Phi^{r,s}_t$, the distribution of bequests received by type (z^c_r, v^c_s), is determined in equilibrium. If an offspring of type (z^c_r, v^c_s) receives a bequest, the distribution $\Phi^{r,s}$ from which the bequest is drawn is given by

$$\Phi^{r,s}_t = \sum_{m,n} \pi(z^c_r, v^c_s \mid z^p_n, v^p_m)\phi^{n,m}_t, \tag{29}$$

where $\phi^{n,m}_t$ is the probability (density) distribution of positive bequests left at time t by households of type (z_n, v_m). For any level of bequests $b > 0$, this density is given by

$$\phi^{n,m}_t(b) = \frac{\sum_j \int I_{x:\{a'_{t-1}(x)=b\}}(1-\psi^{z_n,v_m}_j)dF_{t-1,j}(x : z = z_n, v = v_m)}{\sum_j \int (1-\psi^{z_n,v_m}_j)dF_{t-1,j}(x : z = z_n, v = v_m)}, \tag{30}$$

where the denominator is the total measure of type (z_n, v_m) households deceased at time t and the numerator is the total number of type (z_n, v_m) households deceased at time t leaving a bequest of b.

We denote the proportion of agents passing away between period $t-1$ and t and leaving positive bequests among the total population in period $t-1$ as $\chi_t = \sum_j \int I_{x:\{a'_{t-1}(x)>0\}}(1-\psi^{z_n,v_m}_j)dF_{t-1,j}$. The stationarity of our model implies that this is a constant in the steady state and that the overall probability of receiving bequests needs to be equal to χ_t. Also,

[26] De Nardi (2004), De Nardi and Yang (2016), Fuster et al. (2007), Fuster et al. (2008), among others, explicitly model the intergenerational links.

denote the proportion of agents of type (z_n, v_m) among the deceased who have left positive bequests by $\mu^{n,m}$. This corresponds to the proportion of agents of type (z_n, v_m) leaving positive bequests, which is given by

$$\mu_t^{n,m} = \frac{\sum_j \int I_{x:\{a'_{t-1}(x)>0\}}(1-\psi_j^{z_n,v_m})dF_{t-1,j}(x : z = z_n, v = v_m)}{\chi_t}. \tag{31}$$

Moreover, denote the proportion of children of type (z_r^c, v_s^c) receiving bequests by $\mu_{c,t}^{r,s}$. This is given by

$$\mu_{c,t}^{r,s} = \sum_{n,m} \mu_t^{n,m} \pi(z_r^c, v_s^c \mid z_n^p, v_m^p). \tag{32}$$

By construction, the measure of total households leaving bequests and the measure of those receiving bequests have to be equal.

We are now ready to show how $\Gamma_{j,t}^{r,s}$, the probability of receiving bequests at age j, conditional on being type of (z_r, v_s), is determined in equilibrium. Denote the age distribution of bequest recipients among type (z_r^c, v_s^c) by $P_j^{r,s}$. Then, it follows that

$$\Gamma_{j,t}^{r,s} = \frac{P_j^{r,s} \mu_{c,t}^{r,s} \chi_t}{1+n}. \tag{33}$$

Note that χ_t is the proportion of households leaving bequests in period t within the population in period $t-1$. However, because of population growth, the proportion of households receiving bequests in period t (among the population in period t) is given by $\chi_t/(1+n)$. As we show in Section B.2, we will obtain $P_j^{r,s}$ directly from the data.

A.2 Definition of equilibrium

Given the initial capital stock K_0, an initial measure of households $\{F_{0,j}\}_{j=1}^J$, which implies past average lifetime earnings $\{\mathcal{E}_t\}_{t=-(J-J^E)}^{-1}$, and a sequence of policy parameters (average replacement rate, progressivity of the social security system, and income tax progressivity) $\{\bar{\lambda}_{SS,t}, \lambda_{SS,t}, \lambda_{I,t}\}_{t=0}^{\infty}$, a competitive equilibrium is a sequence of value and policy functions, $\{V_t, V_t^R, \tilde{V}_t, c_t, a'_t, l_t, j_t^R\}_{t=0}^{\infty}$, production plans for firms $\{K_t, N_t, N_{z,t}, Y_t\}_{t=0}^{\infty}$, prices $\{w_{z,t}, r_t\}_{t=0}^{\infty}$, social security variables $\{\tau_{SS,t}, \mathcal{E}_t\}_{t=0}^{\infty}$, income tax variables $\{\bar{\lambda}_{I,t}, \mathcal{I}_t\}_{t=0}^{\infty}$, government spending and debt $\{G_t, D_t\}_{t=0}^{\infty}$, aggregate consumption and wealth $\{C_t, A_t\}_{t=0}^{\infty}$, and population measures $\{F_{t,j}\}_{t=1}^{\infty}$, such that the following statements hold for all $t \geq 0$:

1. The aggregate effective labor supply $\{N_t\}$ is determined by (16).

2. Factor prices $\{w_{z,t}, r_t\}$ are determined competitively from (18) and (19).

3. The value functions $\{V_t, V_t^R, \tilde{V}_t\}$ solve the agent's optimization problem, and the associated policy functions are $\{c_t, a_t', l_t, j_t^R\}$.

4. Given the agent's policy functions, the next period's distribution of agents is given by $F_{t+1,j+1} = \mathcal{F}_{t,j}(F_{t,j})$, where $\mathcal{F}_{t,j}$ is the law of motion induced by the Markov process for labor productivity, the law of motion for skill acquisition, the law of motion for the average lifetime earnings, the endogenous asset accumulation and early retirement decisions and the bequest distribution rule across generations.

5. The average social security taxable earnings follow from

$$\mathcal{E}_t = \sum_j \int_{\boldsymbol{x}:j^R \neq 0} \bar{e} dF_{t,j} / \sum_j \int_{\boldsymbol{x}:j^R \neq 0} dF_{t,j}.$$

6. Average taxable income $\{\mathcal{I}_t\}$, total wealth $\{A_t\}$, and total consumption $\{C_t\}$ read

$$\mathcal{I}_t = \sum_j \int \iota_t(\boldsymbol{x}) dF_{t,j}, \quad (34)$$

$$A_{t+1} = \sum_j \int a_t'(\boldsymbol{x}) dF_{t,j}, \quad (35)$$

$$C_t = \sum_j \int c_t(\boldsymbol{x}) dF_{t,j}. \quad (36)$$

7. The social security tax $\{\tau_{SS,t}\}$ balances the social security budget (20).

8. The income tax level $\{\bar{\lambda}_{I,t}\}$ balances the government budget (21).

9. The government debt and spending $\{D_t, G_t\}$ are constant: $D_t = \bar{D}$ and $G_t = \bar{G}$.

10. The distribution of positive bequests is determined by (29)–(30) and the probability of receiving bequests is determined by equations (31)–(33).

11. The capital market, labor markets for each skill, and the goods market clear:

$$A_t = K_t + D_t, \quad (37)$$

$$N_{u,t} = \sum_j \int_{\boldsymbol{x}:\{j^R=0 \text{ and } z=u\}} yvhl_t(\boldsymbol{x}) dF_{t,j} \text{ for } u \in \{H, L\}, \quad (38)$$

$$Y_t = C_t + G_t + (1+n)K_{t+1} - (1-\delta)K_t.$$

A.3 Computational approach

To solve for the unconstrained and constrained optimal policies, we proceed as follows. First, we create feasible (and reasonable) intervals for each policy instrument ($\lambda_I \in [0.0, 0.3]$, $\lambda_{SS} \in [0.4, 2.1]$, and $\bar{\lambda}_{SS} \in [0.0, 0.80]$) and then discretize each interval using six grid points in each dimension, since the model solution is computationally very intense. In fact, with our choice of grid size, we must compute $6^3 = 216$ transitions for each experiment. Besides, we need to make sure that the social security budget constraint (20) is satisfied for *each* combination of parameters on the grid. High values of average replacement rate $\bar{\lambda}_{SS}$ and high or low values of pension progressivity λ_{SS} can create large financing needs.[27] These considerations determined our choice of the upper and lower limits of the social security parameters on the grid.

Second, for each policy $\mathbf{\Lambda}$ from the grid, we compute the transitional dynamics of the economy from the calibrated steady state with the initial policy $\mathbf{\Lambda}^0$ to a new steady state associated with $\mathbf{\Lambda}$ and obtain the value functions for the currently alive agents, $V_0(\boldsymbol{x}; \mathbf{\Lambda}^0, \mathbf{\Lambda})$, as well as the sequence of value functions for future generations, $\{V_s(\boldsymbol{x}^N; \mathbf{\Lambda}^0, \mathbf{\Lambda})\}_{s=1}^T$, where $\boldsymbol{x}^N$ is the newborn's state and T is the total number of transition periods. Finally, we apply a standard optimization routine in that smaller space that interpolates $V_0(\boldsymbol{x}; \mathbf{\Lambda}^0, \mathbf{\Lambda})$ for each state $\boldsymbol{x}$ and $\{V_s(\boldsymbol{x}^N; \mathbf{\Lambda}^0, \mathbf{\Lambda})\}_{s=1}^T$ for each state $\boldsymbol{x}^N$ and time s using cubic splines at the given candidate solution $\mathbf{\Lambda}$, which does not necessarily lie on the grid.

Our approach is also well-suited to impose "Pareto constraints" such as (24) or (25). Since these constraints are composed only of the value functions of current and/or future generations, we can also express them through cubic spline interpolation using the value functions on the grid. This approach can be applied to a wide range of computationally intensive optimization problems.

B Calibration

B.1 Type-specific survival probability rates

The conditional mortality rates in the model, $1 - \psi_j^{z,v}$, depend on the agent's age j, education level $z \in \{z_L, z_H\}$ and productivity fixed effect $v \in \{v_L, v_H\}$. We assume that they

[27]As it is clear from Figure 1, if the social security system is very progressive (λ_{SS} is much above 1) the pension entitlements of low-income agents are very high. If the social security system is very regressive (λ_{SS} is much below 1) the pension entitlements of high-income agents are very high. In the case of high $\bar{\lambda}_{SS}$, pension entitlements of all agents are very high.

follow a Gompertz force of mortality (i.e., mortality hazard) function

$$\mathcal{M}_j^{z,v} = \mu_1^{z,v} \times \frac{e^{\mu_2(j-1)} - 1}{\mu_2}. \tag{39}$$

The first parameter $\mu_1^{z,v}$ controls how mortality varies with the agent's type (z, v). The second parameter μ_2 controls how mortality changes with age j, conditional on the agent's type (z, v). Given this specification, the (unconditional) probability that the type (z, v) agent of age 1 survives up to age $j > 1$ is given by

$$\mathcal{P}_j^{z,v} = \Pi_{i=1}^{j-1} \psi_i^{z,v} = \exp(-\mathcal{M}_j^{z,v}). \tag{40}$$

The conditional survival probability rate $\psi_j^{z,v}$ can then be computed recursively as:

$$\psi_j^{z,v} = \mathcal{P}_{j+1}^{z,v} / \mathcal{P}_j^{z,v} \text{ with } \mathcal{P}_1^{z,v} \equiv 1. \tag{41}$$

The parameter vector $\boldsymbol{\mu} \equiv (\mu_1^{z_L,v_L}, \mu_1^{z_H,v_L}, \mu_1^{z_L,v_H}, \mu_1^{z_H,v_H}, \mu_2)$ has to be estimated inside the model, since we rely on income-specific mortality rates. To do this, we use the following two independent sources of data. The first is Bell et al. (1992), which computes the age profile of conditional survival probability rates for female and male workers covered by social security. We take their estimates for 2010 and average the reported estimates for both genders using a (constant) share of female workers equal to 50%, consistent with our CPS sample. The second source is Chetty et al. (2016).[28] The authors report conditional mortality rates by gender, age, year, and household income percentile. Household income percentiles are calculated separately for each gender, age, and year. Since mortality rates are not reported for income percentiles defined on a pooled sample of females and males, we use their estimates for the male population only and pool the data for the years 2001–2014. In the calibration, we target mortality rates at age 70, but we also report mortality rates at other ages that were not targeted. To compute the model counterpart of the mortality rates, we use the pre-reform stationary distribution of pre-government income, where pre-government income is equal to $r_t a + e$ for workers and $r_t a$ for retired agents. To ensure consistency between the two data sources, we normalize the income-dependent mortality rates in Chetty et al. (2016), such that the average mortality rate at age 70 coincides with the corresponding estimate in Bell et al. (1992). The obtained estimates are $\mu_1^{z_L,v_L} = 0.0175 \times 10^{-2}$, $\mu_1^{z_H,v_L} = 0.0117 \times 10^{-2}$, $\mu_1^{z_L,v_H} = 0.0078 \times 10^{-2}$, $\mu_1^{z_H,v_H} = 0.059 \times 10^{-2}$, and $\mu_2 = 11.93 \times 10^{-2}$.

[28]The data are publicly available in Table 15 at https://healthinequality.org/data/.

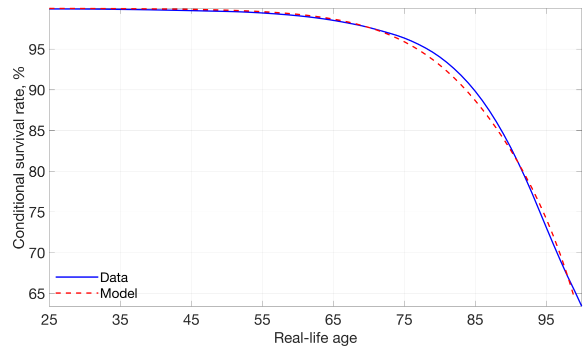

Figure 16: Age profile of conditional survival probability rates in the model and data.
Notes: The empirical rates (dashed line) were constructed based on the data by Bell et al. (1992) for 2010.

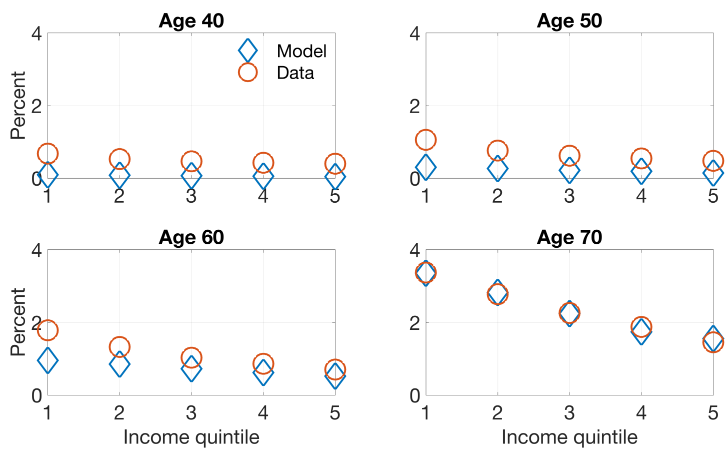

Figure 17: Conditional mortality rates by age in the model and data.
Notes: The empirical rates (dashed line) were constructed based on the data by Chetty et al. (2016) for male individuals. The profile at age 70 was targeted during the calibration.

Figure 16 shows the model fit for the age profile of conditional survival probability rates. Figure 17 (last panel) shows the fit of mortality rates by income quintiles at age 70, which were targeted during the estimation. The remaining panels show how the model fits the mortality rates by income at other ages. These data moments were not targeted.

B.2 Bequests

We start with the intergenerational transition matrix for types $\pi(z^c, v^c \mid z^p, v^p)$. First, we assume that the two processes are independent, namely, $\pi(z^c, v^c \mid z^p, v^p) = \pi_z(z^c \mid z^p)\pi_v(v^c \mid v^p)$. For the diagonal elements of the intergenerational transition matrix for education $\pi_z(z^c \mid z^p)$, we use OECD data from 2012, according to which 64% of parents with at least one completed high-school graduate education have a child with the same educational

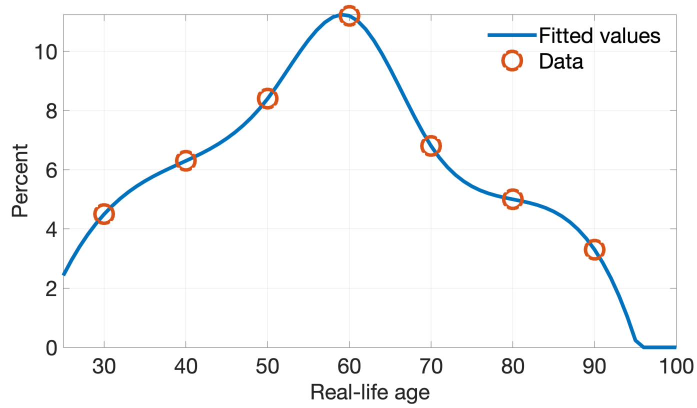

Figure 18: Age profile of probabilities to receive a bequest in the data.

Notes: See Footnote 30 for the data source.

level, while the same statistic for parents with at least one college degree is 61%.[29] The transition matrix for the fixed effect is assumed to be i.i.d. with $\pi_v(v^c \mid v^p) = 0.5$.

Next, we turn to the age- and type-specific probabilities $\Gamma_j^{z,v}$ of receiving a positive bequest in the steady state. Recall that these are determined by equation (33). First, we assume that the age distribution of bequest recipients $P_j^{r,s}$ is independent across types and we denote it by P_j. Second, we take the empirical age profile for the probability of receiving a bequest and fit a third-degree polynomial. Figure 18 shows the original data (circles) and fitted values (solid line).[30] Renormalizing this data by the average probability of receiving bequests gives us P_j, and $\Gamma_j^{z,v}$ can then be calculated from equation (33).

Figure 19 shows the distribution of bequests (normalized by the economy-wide median pre-government income) in the data and model. As reflected by the figure, around 20% of households leave no bequests, a moment that we target in our calibration.

B.3 Social security

According to the social security legislation, an individual's pension benefit depends on their average monthly earnings. The monthly pension benefit of a worker who retires at the normal retirement age is determined by a statutory replacement rate schedule, which is a function of the individual's average monthly earnings. The schedule consists of three brackets with a constant marginal replacement rate of 90%, 32%, and 15% in the lowest, intermediate, and highest brackets, respectively. To bring the empirical schedule of re-

[29] See *Intergenerational mobility in education* at https://stats.oecd.org/.

[30] Data source: https://budgetmodel.wharton.upenn.edu/issues/2021/7/16/inheritances-by-age-and-income-group. Since agents in the model can receive multiple bequests, we consider inheritances from all sources in the data (parents, grandparents, and others).

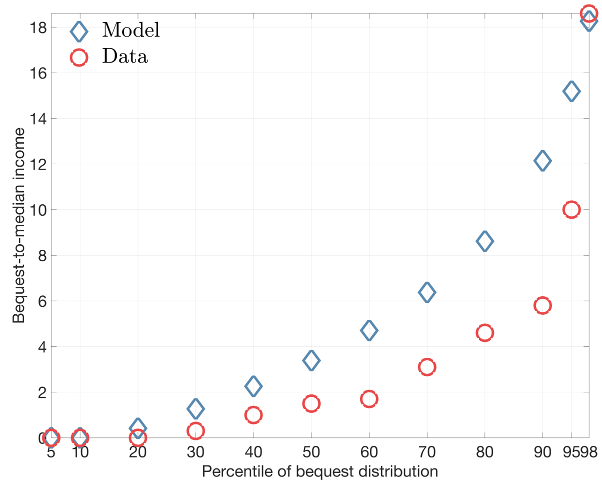

Figure 19: Bequest distribution in the model and data.

Notes: Notes: Bequests are normalized by the economy-wide median pre-government income. The empirical data comes Hurd and Smith (2001).

placement rates to the model, several transformations are required.

First, the worker's average monthly earnings and the brackets are converted into their annual counterparts (multiplying each by 12) since one period in the model is equal to one year. Figure 20 shows the implied schedule of average replacement rates obtained after this step. The two vertical dashed lines correspond to the annualized bend points of $10,620 and $64,032.[31,32]

Second, the bend points are adjusted to the average number of earners in a household since the observation unit in our model is a household, while the statutory schedule above applies to individuals. In the 2018 CPS extract, less than 3% of households consist of 4 or more earners. These households are disregarded in the calculations. In the remaining sample, there are 47% of single-earner, 46% of two-earner, and 7% of three-earner households. Adjusting the first annualized bend point of 10,620 to the number of earners, we obtain: 10,620 × $[1 \times 0.47 + 2 \times 0.46 + 3 \times 0.07]$ = $16,992. We proceed similarly with the second annualized bend point of 64,032 to get $102,451. Third, the household's average annual earnings and the bend points are normalized by the economy-wide average taxable earnings, which we compute from the CPS.

Finally, the parameters $(\lambda_{SS}, \bar{\lambda}_{SS})$ are estimated by fitting the replacement rate function $R(\cdot)$ in (5) to the empirical schedule of average replacement rates depicted in the figure.

[31] All dollar amounts in this subsection are given in 2017 U.S. dollars.

[32] The parameters of the statutory pension benefit formula can be found in Table 2.A11 in the Supplement to the annual Social Security report available at `https://www.ssa.gov/policy/docs/statcomps/supplement/2017/index.html`.

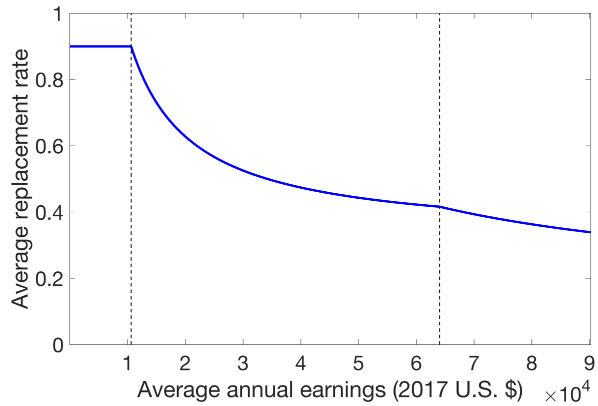

Figure 20: Statutory replacement rate schedule.

Notes: The figure visualizes the statutory relationship between the individual's average earnings and their average replacement rate. The figure is constructed based on the statutory pension benefit formula from 2017 that applies to all individuals who retired at the normal retirement age. In the data, the formula is expressed in monthly terms, while the figure shows the annualized values. The dashed lines correspond to the bend points equal to $10,620$ and $64,032$ (in 2011 U.S. dollars). The statutory marginal replacement rates in the lowest, intermediate, and highest brackets are 90%, 32%, and 15%, respectively.

The estimated value of the replacement rate level $\bar{\lambda}_{SS}$ is then adjusted inside the model, keeping progressivity λ_{SS} fixed, such that the social security tax, τ_{SS}, matches exactly its empirical counterpart.

During the calibration, the cap is estimated inside the model by targeting the share of households in the CPS whose total pre-tax earnings exceed the maximum taxable earnings threshold adjusted for the number of earners. The statutory value of the cap is $127,200$ in 2017.[33] Adjusting this value for the number of earners in a household computed above, we obtain 127,200 × $[1 \times 0.47 + 2 \times 0.46 + 3 \times 0.07]$ = 203,520. Hence, cap is calibrated to target the share of households with earnings above this threshold.

[33]The data on the maximum taxable earnings threshold are taken from Table 2.A3 (see the previous footnote for the data source).

C Additional Figures

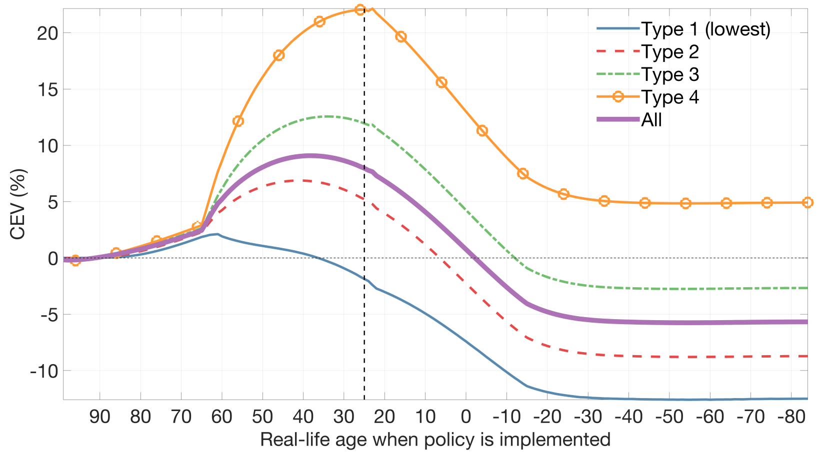

Figure 21: Welfare effects by type and age in the full optimal reform (CEV, %).
Notes: See notes of Figure 8.

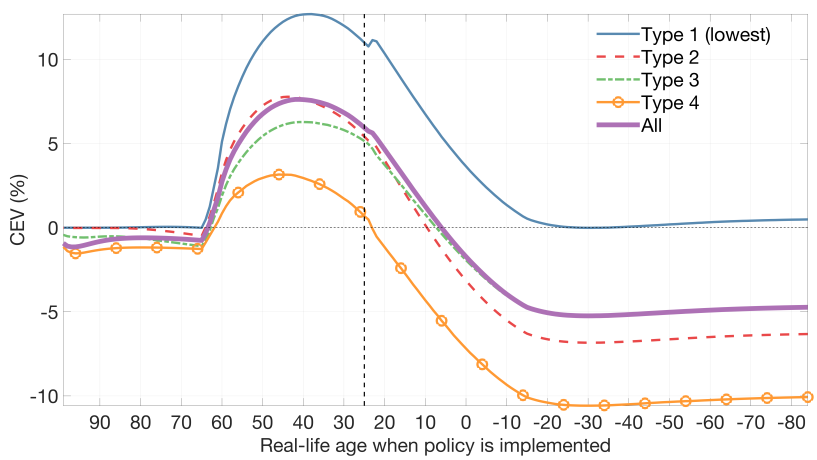

Figure 22: Welfare effects by type and age under the *"Alternative pension"* reform.
Notes: The reform maximizes the welfare of all currently alive generations subject to the constraint that no Type 1 agents – whether currently alive, along the transition path, or in the final steady state – are made worse off (on average). See also notes to Figure 8.

D Robustness

D.1 Results with redistributional constraints

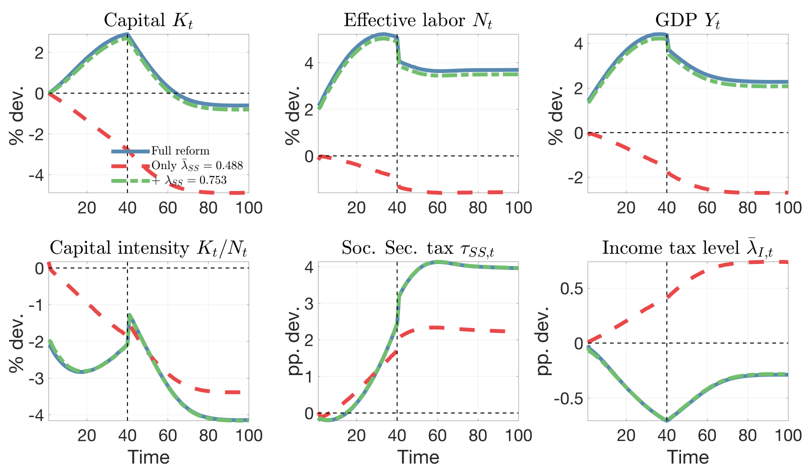

Figure 23: Decomposition of transitional dynamics under *No cohort is worse off* reforms.

Notes: The reform maximizes the welfare of all currently alive generations subject to the constraint that no cohort is made worse off (on average). See also notes to Figure 14.

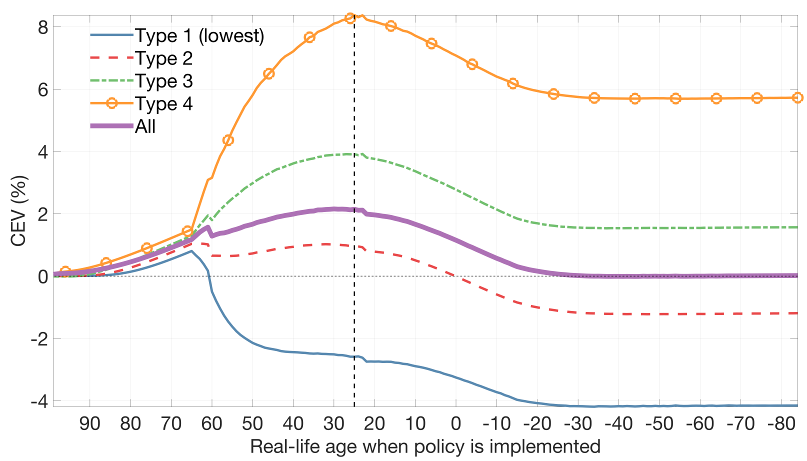

Figure 24: Welfare effects by type and age under the *"No cohort is worse off"* reform.

Notes: The reform maximizes the welfare of all currently alive generations subject to the constraint that no cohort – whether currently alive, along the transition path, or in the final steady state – is made worse off (on average). See also notes to Figure 8.

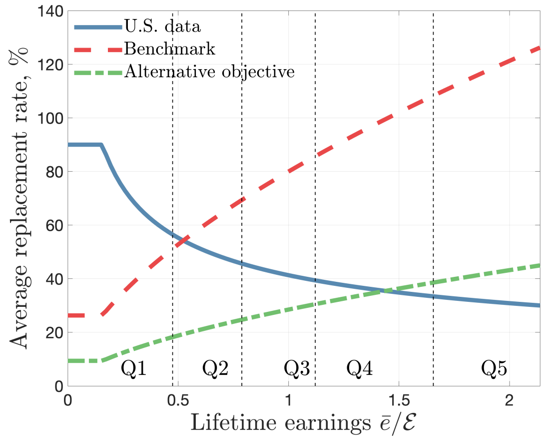

Figure 25: Replacement rates in the unconstrained alternative and benchmark cases.

Notes: The dashed vertical lines mark the quintiles of the lifetime earnings distribution among retired agents in the pre-reform stationary equilibrium. The average lifetime earnings $\bar{e}$ in the model and data are normalized by average taxable earnings $\mathcal{E}$.

D.2 Results with alternative objective

In this appendix, we present the results with a social welfare function that maximizes the welfare of households that are born into the final steady state, while potentially imposing the Pareto constraints defined in Section 3. This alternative welfare function is given by:

$$W^A(\boldsymbol{\Lambda}) = \int V_\infty(\boldsymbol{x}; \boldsymbol{\Lambda}) dF_{\infty,j=1}, \tag{42}$$

where $V_\infty(\boldsymbol{x}; \boldsymbol{\Lambda})$ is the value function of a newborn in the final steady state.[34] In this reform, the planner solves:

$$\boldsymbol{\Lambda}^\star = \arg\max_{\boldsymbol{\Lambda}} W^A(\boldsymbol{\Lambda}). \tag{43}$$

The results of the unconstrained reform are displayed in Table 5 in the main text, and they are discussed in Section 6.2. Figure 25 in this appendix shows how the schedule of average replacement rates changes in this alternative case. Compared to the benchmark case, the pension system also becomes regressive, but pensions now decrease with respect to the status quo for almost all households except for the top quintile, with the most pronounced drop occurring in the lowest two quintiles. As a result, the higher implicit taxes in this reform tend to generate a drop in labor supply.

However, as shown in Figure 26 (Full reform), the effect of higher implicit taxes is offset by the big drop in both the income and social security taxes in the full reform, which

[34]As opposed to the benchmark, the individual value functions and the social welfare function do not depend on $\boldsymbol{\Lambda}^0$ in this case.

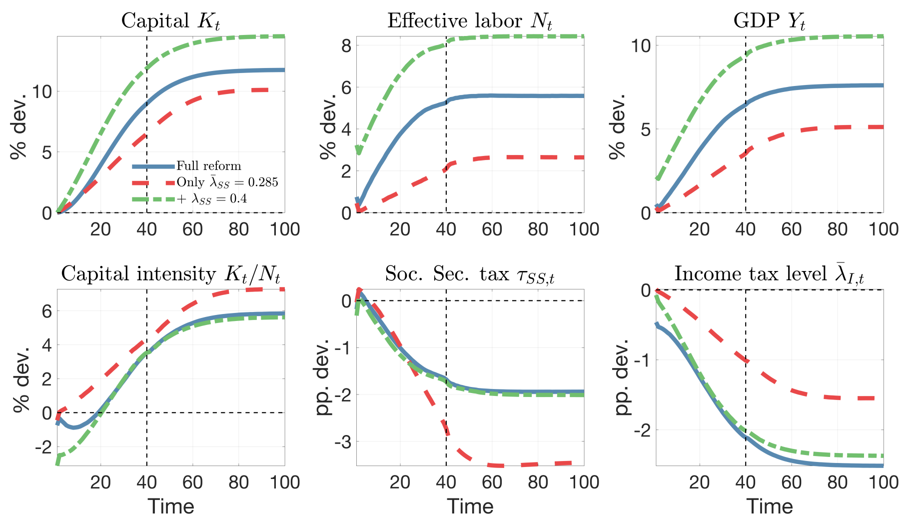

Figure 26: Transitional dynamics with alternative social welfare function.

Notes: The reform maximizes the welfare of newborn generations in the final steady state. See also notes to Figure 14.

is possible due to the reduction in pension generosity. This forces most households to increase their savings and labor supply for retirement. In turn, the relatively big increase in capital accumulation translates into an increase in the capital-labor ratio and, hence, wages, further reinforcing the increased incentives for aggregate labor supply. Note also that the increases in both labor supply and aggregate capital lead to a significant rise in the long-term output level, generating big welfare gains for future generations at the expense of current ones, as they face a less generous pension system but have limited benefits from the long-run efficiency gains.

Figure 26 also decomposes the transitional paths for each policy instrument in this reform. As in the benchmark case, pension progressivity again has a very pronounced impact on both capital accumulation and labor supply, highlighting the power of this instrument in reducing aggregate distortions. Figure 27 illustrates another similarity with the benchmark reform. The alternative reform generates again a strong pattern of redistribution towards high-productivity (rich) households for all generations, as Type 4 households consistently gain, while the income-poorest agents (Type 1) mostly lose. As a result, within-cohort inequality increases. Note, however, that this reform prescribes an increase in income tax progressivity, which imposes additional distortions, to partially offset the increased inequality created by the change in the pension system.

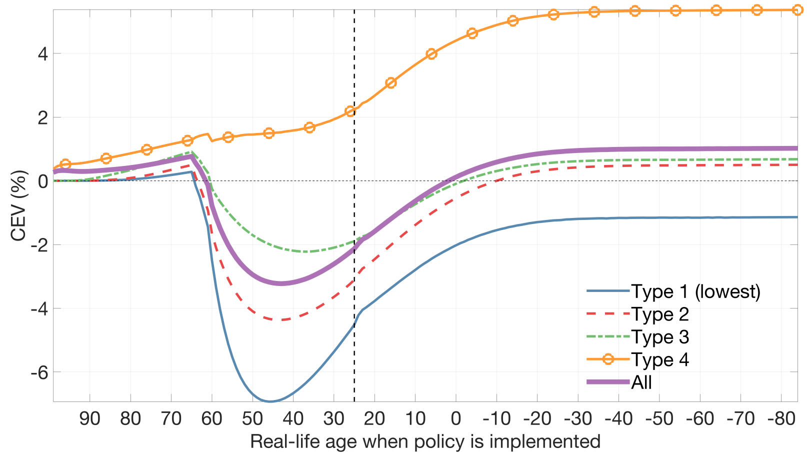

Figure 27: Welfare effects in the reform with alternative social welfare function (CEV, %).
Notes: The reform maximizes the welfare of newborn generations in the final steady state. See also notes to Figure 8.

D.3 Results from counterfactual experiments

Table 6: Optimal policy in counterfactual experiments in the unconstrained reform.

	Policy			CEV, %		Support,
	$\bar{\lambda}_{SS}$	λ_{SS}	λ_I	Alive	Future	%
Status quo	0.413	1.420	0.216	–	–	–
Benchmark	0.800	0.400	0.191	7.929	-5.673	89.3
Exogenous human capital	0.800	2.100	0.245	11.779	0.711	67.5
Exogenous retirement	0.800	2.100	0.116	7.505	-4.846	82.9
Exogenous capital income taxation	0.800	2.100	0.122	7.710	-5.990	83.6
Partial equilibrium	0.800	2.100	0.102	7.268	-3.501	77.4
No taxation of pension benefits	0.700	0.400	0.240	7.825	-8.446	95.7
No deduction of SS contributions	0.800	0.400	0.191	8.301	-5.769	89.7
Same mortality rates	0.800	0.400	0.186	7.758	-5.112	93.3

Notes: In all experiments, the government maximizes the welfare of currently alive generations subject to no constraints on redistribution. See notes to Table 2.

Table 7: Counterfactual experiments with alternative pension policy.

	Policy $\bar{\lambda}_{SS}$	λ_{SS}	λ_I	CEV, % Alive	Relative gain, %
Status quo	0.413	1.420	0.216	–	–
Exogenous human capital	0.800	0.400	0.191	4.190	35.6
Exogenous retirement	0.800	0.400	0.191	7.190	95.8
Exogenous capital income taxation	0.800	0.400	0.191	7.062	91.6
Partial equilibrium	0.800	0.400	0.191	7.006	96.4
No taxation of pension benefits	0.717	2.100	0.165	5.270	67.3
No deduction of SS contributions	0.717	2.100	0.165	5.887	70.9
Same mortality rates	0.717	2.100	0.165	7.612	98.1

Notes: For the first four experiments, we impose the unconstrained optimal policy from the benchmark case. For the last three reforms, we impose the optimal policy from the *Alternative pension reform*. The last column shows the welfare gain under the exogenous policy as a percentage of the gain under the respective optimal policy from Table 6.