Optimal Fiscal Reform with Many Taxes[1]

Daniel Carroll

Federal Reserve Bank of Cleveland
daniel.carroll@clev.frb.org

André Victor D. Luduvice

Federal Reserve Bank of Cleveland
andrevictor.luduvice@clev.frb.org

Eric R. Young

University of Virginia
ey2d@virginia.edu

Abstract

We study the optimal one-shot tax reform in the standard incomplete markets model where households differ in their wealth, earnings, permanent labor skill, and age. The government can provide transfers by raising tax revenue and has several tax instruments at its disposal: a flat capital income tax, a flat consumption tax, and a non-linear labor income tax. We compute the equilibrium and transitional dynamics for 3888 different tax combinations and find that the optimal fiscal policy funds a transfer that is above 60 percent of GDP through a combination of very high taxes on consumption and capital income. The labor tax schedule has a high average rate and more progressivity than the current US system. We explore the role of transitional dynamics, debt issuance, intergenerational disagreement, and fiscal spending rules in shaping the optimal policy. Policy is broadly similar if it is determined through majority voting rather than by a utilitarian planner.

KEYWORDS: Optimal Taxation; Inequality; Heterogeneous Agents; Incomplete Markets; Voting;
JEL CLASSIFICATION CODES: E62, E21, D72, H21

[1]This paper has benefited from the computational resources provided by the BigTex High Performance Computing Group at the Federal Reserve Bank of Dallas. We would like to thank Anmol Bhandari, Ayşe İmrohoroğlu, Chengcheng Jia, Barış Kaymak, Dirk Krueger, Jonathan Heathcote, Roozbeh Hoseini, Sewon Hur, Kurt Lunsford, Gastón Navarro, Tom Phelan, Kjetil Storesletten, Murat Tasci, Gustavo Ventura, Nicolas Werquin, as well as the participants at the Cleveland Fed Brown Bag, the Dallas Fed Brown Bag, the SEA 2022, the BSE Summer Forum 2023, the SED 2023 and 2024, the IIPF 2023, the USC Annual Macro Day 2023, the LubraMacro 2023, the 2024 System Equitable Growth Conference, and the EPGE-FGV/RJ seminar for helpful comments. We also thank Cornelius Johnson for outstanding research assistance and Stephanie Tulley for outstanding data and reference support. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Cleveland or the Federal Reserve System. This version: March 2025. First draft: October 2022.

1 Introduction

The past half-century has been marked by an increase in earnings and wealth inequality. This trend has raised questions about whether the government should redistribute income to reduce economic disparities and, if so, how it should achieve this aim. Which revenue sources should be taxed? At what rates? How progressive should taxation be? How much income should be redistributed? These questions are especially challenging when the revenue for financing redistribution comes primarily from distortionary taxation. More equal divisions of the economic pie necessarily come at the cost of reducing the pie itself, creating the classic "equity/efficiency" trade-off.

To understand how different types of taxes and their resulting amount of redistribution shape the optimal fiscal policy, we turn to the workhorse model of inequality in macroeconomics: the dynamic general equilibrium model with incomplete asset markets and binding borrowing constraints (Bewley, 1986; Imrohoroğlu, 1989; Huggett, 1993; Aiyagari, 1994). In the standard model, infinitely lived agents face uninsurable idiosyncratic labor productivity risk which generates income and wealth inequality across households. We enrich the model by adding a stochastic aging structure and augmenting the productivity process with permanent labor skill differences and with high-productivity states. These changes allow the model to incorporate overlapping generations with a retirement phase while endogenously generating the upper tails in earnings and wealth that are present in the data. These tails are critical to our question since they are simultaneously an attractive target for redistribution and an elastically supplied input to future production.

We then use the model as a quantitative laboratory to study how permanent reforms to the current tax/transfer system affect both aggregate activity and household welfare. We evaluate the effects from a wide range of fiscal policy tools, including consumption taxes, capital income taxes, and labor income taxes, as well as the progressivity of the labor income tax schedule. The government uses tax revenue to finance government expenditures, service existing debt, and provide uniform lump-sum transfers.

Starting from the calibrated steady state, we solve for the fiscal policy reform that maximizes utilitarian social welfare, taking into account the transitional dynamics to the new steady state. The optimal policy features a very high level of taxation. Within a finite, but wide and numerous, set of tax policy combinations, the optimal choice of tax rates for

consumption and capital income are 51.2 and 60.0 percent, respectively, while the average labor tax rate is 57.0 percent. These very high levels of taxation support large lump-sum transfers roughly equal to 60.0 percent of GDP. At the same time, the optimal policy has higher progressivity than under the benchmark calibration, ensuring that top earners bear the brunt of the high labor income tax.

This large and comprehensive tax bill, coupled with a massive windfall redistribution, greatly depresses economic activity. In the long run under the optimal schedule, aggregate capital and average hours worked plummet by more than 40 percent. Aggregate output falls by nearly one-third, while transfers rise to 17 times their initial level.

Despite severely reducing aggregate activity in the model economy, the optimal fiscal policy generates much higher average welfare. Two features of the environment are responsible for this result: a high degree of inequality in the initial steady state and very limited means for households to insure themselves against large and persistent idiosyncratic income shocks. These conditions ensure that there is both a strong demand for social insurance and a ready supply of revenue. Support for massive redistribution is so widespread in the model that efficiency considerations have little weight in moderating the policy outcome. While it is the case that introducing additional features to increase the tax elasticity of revenue and pull the peak of the Laffer curve to the left could constrain redistribution in equilibrium, the underlying support for a very high level of taxation would still hold so long as high levels of idiosyncratic risk remain present.

The optimality of this policy arises from an assignment of higher weights to the welfare of low-income households that is implicit in a utilitarian social planner's objective. These households' income derives almost entirely from their labor, which makes taxing capital attractive to them, and since the equilibrium policy is progressive, a large transfer can be secured while keeping their labor income tax burden to a minimum. Although low-income households are still targeted by the consumption tax, in absolute terms their consumption levels are low; so the transfer achieved under the optimal policy more than offsets any consumption taxes paid. If the wealth distribution is appropriately calibrated to the US data, low-income households are both numerous and have high marginal utilities of consumption; so their welfare has disproportionate weight in the planner's objective.

We define a grid for each tax instrument and then construct a tax menu with 3888 policies consisting of all combinations over these grids. The optimal tax structure pushes all

tax parameters to their maximal values with the exception of labor tax progressivity.[2] The equilibrium policy essentially strikes a balance between two effects from greater progressivity: increasing overall welfare by shifting the tax burden away from low- and middle-income earners and reducing output (and as a consequence transfers) by discouraging the most productive households from working. Progressivity is also the dimension along which wealth-poor households disagree. Those with high earnings support a labor tax schedule with a moderate rate on average and a high degree of progressivity. In contrast, the lowest earners prefer the opposite: a high flat labor tax so as to maximize the transfer. Those in the middle also favor a high labor tax rate but would give up some of the transfer for more progressivity in order to keep their own tax bills low. Because in equilibrium revenue declines as progressivity increases, this tradeoff also appears in the optimal policy. Intuitively, the solution to the planner's problem can be characterized in two steps: first, increase the average level of taxation to fund a very high transfer; second, trade back some of the revenue to buy progressivity and minimize the tax burden of the poor.

The optimal policy is not simply a by-product of the unequal initial distribution of capital. Repeating the optimal policy exercise starting from a much starker low-wealth distribution – specifically, the one that arises in the long run under the baseline's high-tax optimal policy – delivers nearly the same result. The only difference between the baseline and the low-wealth solutions is that progressivity is higher in the second case because consumption is already distributed much more equally in the low-capital environment. Another reason for the robustness of the high-tax policy across distributions is that the costs of transitioning back to a low-tax economy with greater aggregate activity are so high that it is welfare maximizing to keep the status quo. Households, in this sense, have to abide by the policies of the past. If policy is instead chosen to maximize average steady-state utility, meaning that transition costs are ignored, high taxation and very high transfers are still optimal. An important difference in this case, however, is that when transitions are not internalized, it is optimal to eliminate capital income taxation and maximize progressivity.[3]

The critical role of the transitional dynamics for determining the optimal tax schedule, and especially the optimal capital income tax rate, motivates an examination of the preferences of households born in different periods. If the average welfare of future newborns were

[2]The details of the grids are explained in Section 4.

[3]This result is a deviation from Aiyagari and Peled (1995), where optimal long-run capital taxes are positive, and is an indicator of the importance of multiple instruments and progressive tax options.

valued by the social planner rather than only the initial living one, how would the optimal tax choice change? Tracing out each cohort's optimal "from behind the veil of ignorance" policy reveals that the principal inter generational conflict is over the capital income tax. All households born after 15 periods into the transition would prefer that the capital income tax had been eliminated, mirroring the result from the steady-state-only optimization.

The optimal degree of progressivity is constrained not only by Laffer curve considerations for the transfer, but also by the government's spending and debt service commitments. In our baseline economy, the government is committed to sustain the initially calibrated levels of spending and debt throughout the transition. This fiscal regime is especially onerous under the optimal policy as that policy severely depletes the tax base and causes the government's fiscal obligations to grow as a share of GDP. As an alternative, we consider a regime under which the government keeps the *ratios* of government spending and debt to GDP constant. With these new rules in place, the cost of financing the government responds to economic activity, allowing larger transfers to be sustained with the same tax schedule; however, the optimal policy in this less-stringent fiscal regime is quite similar to that in the baseline. The only difference is that labor income taxes are even more progressive when the government's budget constraint is relaxed. Solving a series of intermediate regimes in which only one obligation (either debt or government spending) is fixed in levels while the other is maintained as a constant fraction of GDP illustrates again the idea of buying progressivity. As the fiscal regime becomes less fiscally stringent (in terms of the cost relative to GDP), the optimal policy becomes more progressive while also delivering a higher transfer-to-output ratio.

Taken together, our results show that the widely used model of inequality in macroeconomics, when given enough freedom over the structure of tax policy, predicts that strong support for very large fiscal transfers should be observed in highly unequal societies. However, among the OECD countries, no economy imposes distortionary taxes at the levels suggested by the model. Two potential explanations for this spring to mind. First, the model may not fully capture the efficiency loss from distortionary taxes, leading it to overstate the feasible level of transfers. For instance, productive economic activities like entrepreneurship and human capital accumulation may be especially sensitive to high taxation. Nonetheless, introducing these additional factors seems unlikely to materially alter our findings given low-income households' general preference for redistribution in lieu of economic activity in the

model.[4] Furthermore, we also observe that large reductions in the equilibrium wage, present in the model's outcome under optimal policy, do little to reduce low-income workers' desire for large transfers. Households care about the size of the economic pie only in so much as it affects the size of their slice.

Second, the model may not accurately reflect the process by which equilibrium policy is determined in the data. Utilitarian social planners, while a useful theoretical benchmark, do not exist in the real world. In most advanced economies, a democratic political process ultimately determines the level of taxation and redistribution. Following the strategy in Carroll et al. (2021), we identify fiscal policies that could arise under a majority vote. Because of the multidimensionality of the policy space, majority rule has the potential to deliver more than one equilibrium. While our method allows for such a case, we find that only one policy survives political competition, and this policy is nearly identical to the utilitarian planner's policy. The only difference is that the policy found under majority vote has higher progressivity. Again, the high population share of the poor and their unanimity of preference for redistribution in the model economy would, under a majority vote, decide elections in favor of high taxes. However, this equivalence holds only if each household votes in accordance with its fiscal interest as captured by its individual state space in the model. If, instead, low-income households have less political power or if they place less importance on tax policy than on other non-tax issues, the equilibrium under voting features more moderate taxes and transfers. In a follow-up short paper (Carroll et al., 2024), we further explore this issue and investigate some of the drivers of individuals' preferences over redistribution in the data. We provide empirical evidence that political identity and ideology are strong predictors of people's preferences over taxes and transfers.

Related Literature. Our paper contributes to the quantitative macroeconomics literature that focuses on the positive, normative, and political effects of public finance reforms through the lens of heterogeneous-agent models. Our paper is closely related to the established literature on the effects of optimal income and labor tax progressivity in this class of models, such as in Bakış et al. (2015), Guner et al. (2016), Heathcote et al. (2017), Imrohoroğlu et al. (2018), Holter et al. (2019), and Kindermann and Krueger (2022). We add to this literature

[4] An example of this is that a utilitarian social planner would choose a higher capital income tax rate in an environment with capital-skill complementarity than in one with a Cobb-Douglas production function as shown in Kina et al. (2024).

by expanding the usual focus on optimal progressivity to an environment where all tax rates available in the government's menu are chosen simultaneously along with the curvature of the progressive tax function.

In recent analyses, Dyrda and Pedroni (2023), Acikgoz et al. (2018), Boar and Midrigan (2022), Ferriere et al. (2022), Guner et al. (2023b), Jones and Li (2023), Abraham et al. (2023) and Macnamara et al. (2023) study optimal tax reforms with particular focus on the flexibility of the tax menu available to the government. The first two papers allow labor and capital taxes to be time-varying along the transition, approximating the solution to an unrestricted Ramsey problem. The next two have flexible tax functions that approximate the US tax schedule, and the welfare optimization is conducted over the parameters that govern them. In Ferriere et al. (2022) those parameters govern income progressivity and transfers, while in Boar and Midrigan (2022), they govern both income and wealth taxes, with a lump-sum transfer. Abraham et al. (2023) optimize over a joint objective, mixing the tax-transfer system with Social Security while Macnamara et al. (2023) use a similar set of taxes as ours but focusing on steady-state solutions. Jones and Li (2023) restrict their focus on a flexible set of parameters for the Social Security benefit schedule. Finally, Guner et al. (2023b) optimize over different mixes of taxes, including consumption and wealth taxes, for a revenue-generating objective with a minimal welfare cost.

We add to this literature by allowing for flexibility and simultaneity within a general menu of taxes in a rich model with workers, retirees, and heterogeneous skill levels. In addition, as in Carroll et al. (2021), we track households' value functions under every policy combination and over the whole state-space, enabling us to characterize how a household's state alters their desired tax menu. We also document how the optimal policy changes depending upon different fiscal regimes or on alternative forms of preference aggregation (e.g., head-to-head voting). This approach allows us to document the effect of each additional model feature on the optimal tax menu choice. In particular, the presence of retired households and progressive labor taxes is a critical departure from Carroll et al. (2021) and leads to a very different policy outcome. One of the mechanisms that becomes relevant in our setting is the interaction of consumption taxation with labor tax progressivity, which together work as important sources of tax revenue and redistribution while affecting workers and retirees differently. The relevance of this interaction in heterogeneous-agents models is also present in Conesa et al. (2023) and Macnamara et al. (2023).

Our results on the welfare optimality and political preference for large lump-sum transfers are related to a recent strand of the literature that studies the effects of transfers, often in the form of net income tax and universal basic income reforms, in incomplete markets models. Recent examples include Lopez-Daneri (2016), Conesa et al. (2023), Daruich and Fernandez (2024), Guner et al. (2023a), and Luduvice (2024). Moreover, we find that a progressive income tax schedule dominates an affine one, consistent with Heathcote and Tsujiyama (2021).

We also engage with a branch of the literature in the same methodological tradition that focuses on the weights that a social planner would choose to implement tax schedules and changes observed in the data. Some of the papers that discuss these Pareto weights, also in the context of progressivity choice, are Chang et al. (2018) and Wu (2021). The Downsian voting equilibria we analyze can be thought of as capturing whatever political mechanism led to a given policy by constructing the implicit planning weights that would support that outcome. Our paper is also connected to papers that focus on political equilibria in heterogeneous-agents incomplete markets economies like Aiyagari and Peled (1995), Krusell et al. (1997) Corbae et al. (2009), Bachmann and Bai (2013), and de Souza (2022). In our discussion of alternative aggregation methods beyond the social planner, we employ the Downsian voting approach used in Carroll et al. (2021).

Road Map. The paper is organized as follows. Section 2 constructs the quantitative model and defines the recursive competitive equilibrium. Section 3 describes the calibration used to map the model to the data. Section 4 presents the results of our numerical experiment. Section 5 reports the optimal reform and explores the factors underpinning it. Section 6 extends our findings to a multidimensional voting equilibrium concept. Finally, Section 7 concludes the paper.

2 Model

We build a heterogeneous-agents model with incomplete markets as in Aiyagari (1994), Bewley (1986), Huggett (1993), and Imrohoroğlu (1989), augmented with overlapping generations, exogenous retirement, heterogeneous permanent skill levels, elastic labor supply, and a menu of distortionary taxes, including consumption and progressive labor income tax-

ation. Time is discrete. In any period t there is a unit continuum of households that are either young or old. Each household is endowed with one unit of discretionary time. The fundamental difference between the two household age types is that the young may supply labor with their discretionary time, while old households may only use it to consume leisure. For this reason, we use the descriptors "young" and "worker" interchangeably, and the same is true for "old" and "retiree." Denote the age of a household by $a \in \mathcal{A} \equiv \{W, R\}$. There is a measure μ_s of skilled households and μ_u of unskilled households with $\mu_u = 1 - \mu_s$. A household's permanent skill type is denoted by $j \in \mathcal{J} \equiv \{u, s\}$. Agents are also heterogeneous with respect to their idiosyncratic productivity shock, $\varepsilon \in \mathcal{E}$, with $\mathcal{E}$ finite, and asset holdings $k \in \mathcal{K}$. The state space of the economy is then the set $X \equiv \{\mathcal{K} \times \mathcal{E} \times \mathcal{J} \times \mathcal{A}\}$.

Demographics. Households age stochastically in a perpetual youth life-cycle as in Yaari (1965) and Blanchard (1985). At the end of each period, a fraction ψ_a of young households age and enter the next period as old households. Likewise, a fraction ψ_d of old households die. When a household dies, its assets are transferred to a newborn household of the same skill type. The initial idiosyncratic productivity of this newborn is drawn from the skill-specific invariant distribution of productivity.

Preferences. Households consume non-durable goods, c, choose how much labor to supply, h, and save using a risk-free, non-state-contingent asset, k, which they may borrow up to an exogenous limit, k_b. They maximize their discounted expected lifetime utility represented by

$$\mathbb{E}_{\not\vdash} \left[\sum_{t=1}^{\infty} \beta^t \psi_a \, u\left(c, h\right) \right], \tag{1}$$

where $\beta \in (0, 1)$ is the discount factor and $\mathbb{E}_{\not\vdash}$ is the expectation operator conditioned on the initial state at $t = 0$.

Technology. There is a Cobb-Douglas production technology over aggregate capital and effective labor, $F\left(K, N\right) = ZK^{\alpha}N^{1-\alpha}$ for $\alpha \in (0, 1)$. Capital depreciates at a rate δ. A stand-in firm operates this technology and behaves competitively. The zero-profit conditions

are

$$r_t = \alpha A \left(\frac{K_t}{N_t}\right)^{\alpha-1} - \delta \tag{2}$$

$$w_t = (1-\alpha) A \left(\frac{K_t}{N_t}\right)^{-\alpha}. \tag{3}$$

Workers and the Labor Market. At the beginning of a period, every worker household draws labor productivity, $\varepsilon \in \mathcal{E}$. Following Kindermann and Krueger (2022), we augment the productivity process by including two high-productivity states in order to match the upper tails of the earnings and wealth distributions. Productivity draws then can be either "regular" or "high", with the high productivity states being independent of skill. Regular productivity states are assumed to follow a Markov chain with skill-dependent transition probabilities $\pi_u(\varepsilon, \varepsilon')$ and $\pi_s(\varepsilon, \varepsilon')$ with corresponding invariant distributions, $\overline{\pi}_u(\varepsilon)$ and $\overline{\pi}_s(\varepsilon)$. High productivity states are drawn according to their respective calibrated transition probabilities.

As in Kaymak and Poschke (2016) and Kindermann and Krueger (2022), among others, the "superstars" structure allows us to accurately capture the concentration in the top tails of both the earnings and wealth distributions in addition to broad measures of inequality – measured here by the Gini coefficient. This strategy dates back to Castañeda et al. (2003) and has become a popular modeling choice in the literature.[5] Superstar states are also consistent with the empirical and structural evidence presented in Kaymak et al. (2023) which finds that earnings differences are the primary driver of top wealth concentration in the US and account for between one-half and two-thirds of total income at the top 1 percent of income. These authors emphasize that theories based on earnings concentration are well-suited to accurately model wealth concentration.

The worker's effective labor is $exp(\varepsilon) \cdot h$. For each unit of effective labor supplied, a worker receives a wage $\zeta(j) \cdot w$, where $\zeta(s) > \zeta(u) = 1$ represents the skill premium. Letting y denote a worker's total labor earnings, $y_{j,t}(h, \varepsilon) = \zeta(j) \cdot w_t \cdot \exp \cdot (\varepsilon) \cdot h$.

Government and Taxes. A government collects taxes and issues new debt, B_{t+1}, and spends these revenues on three types of expenditures: government expenditures G_t, non-retirement transfers Υ_t, and debt service $(1 + r_t) B_t$.

[5] See Boar and Midrigan (2022) and Dyrda and Pedroni (2023) for two recent examples.

The government has three tax instruments: a flat tax on consumption, $\tau_{c,t}$, a flat tax on capital income, $\tau_{k,t}$, and a progressive tax on labor earnings, $T_h(y_t)$. Following what is now standard in the literature (Benabou, 2002; Heathcote et al., 2017; Holter et al., 2019), we employ the non-linear tax schedule

$$T_h(y_t) = y_t \cdot \left[1 - (1 - \tau_{y,t})\tilde{y}_t^{-\nu_{y,t}}\right], \tag{4}$$

where $\tau_{y,t}$ controls the average labor taxes of the economy, $\nu_{y,t}$ controls the curvature of the function and hence its degree of progressivity, and $\tilde{y}_t$ is the ratio of the households' gross earnings, $y_{j,t}(h, \varepsilon)$, to the aggregate average labor earnings, AE_t.

In any period t, the government budget constraint is

$$G_t + (1 + r_t) B_t + \Upsilon_t = \tau_{c,t} C_t + TN_t + \tau_{k,t} r_t A_t + B_{t+1} \tag{5}$$

where TN_t is the aggregate level of labor income tax revenue.

Retirement and Social Security. The government also manages the Social Security system that has its budget balanced separately. It disburses aggregate benefits B and funds them with an additional flat tax on earnings, τ_{SS}. Total Social Security contributions are capped at $\bar{t}_{SS}$, which is used to close the system's budget.

When a household ages into retirement, it begins receiving a Social Security benefit $b_j(\varepsilon)$, which is indexed to the household's last labor productivity draw from its working life. We follow the US Social Security schedule to calculate the payments:

$$b_j(\varepsilon) = \begin{cases} r_1 \bar{y}_j(\varepsilon), & \text{if } \bar{y}_j(\varepsilon) \leq b_1 AE \\ r_1 b_1 \bar{y}_j(\varepsilon) + r_2 \left(\bar{y}_j(\varepsilon) - b_1 \bar{y}_j(\varepsilon)\right), & \text{if } b_1 AE < \bar{y}_j(\varepsilon) \leq b_2 AE \\ r_1 b_1 \bar{y}_j(\varepsilon) + r_2 b_2 \bar{y}_j(\varepsilon) + r_3 \left(\bar{y}_j(\varepsilon) - b_2 \bar{y}_j(\varepsilon)\right), & \text{o.w.} \end{cases} \tag{6}$$

where $\bar{y}_j(\varepsilon)$ is the average labor earnings of a household with skill j and labor productivity ε; AE is the aggregate average earnings in the economy; r_1, r_2, r_3 are the replacement rates at the different levels of income; and b_1, b_2 are the function bend points.

2.1 Recursive Household Problem

The individual state space of the households is $x \equiv [k, \varepsilon, j, a] \in X$. For convenience, we henceforth omit the time subscript in the definition of the recursive household problem. In some instances, it will also be useful to distinguish between households of different ages and skills. In those cases, let f_j^a indicate that f, which may be a value function, decision rule, or subset of X, is associated with a household of age a and skill j. For example, the policy function for consumption of a skilled worker is

$$g_c(k, \varepsilon, s, W) \equiv g_{s,c}^W(k, \varepsilon).$$

Taking a tax policy $\boldsymbol{\tau} \equiv \{\tau_c, \tau_k, \tau_y, \nu_y\}$ as given, a household chooses its consumption c, asset holdings k', and, if young, its labor supply h so as to maximize expected lifetime utility. The problem of a retired household with assets, k, productivity type, ε, and skill, j, is

$$\begin{aligned} V_j^R(k, \varepsilon) &= \max_{c,k'} u(c, 0) + (1 - \psi_d)\beta V_j^R(k', \varepsilon) \\ &\text{s.t.} \\ &(1 + \tau_c)c + k' = (1 + (1 - \tau_k)r)k + (1 - \tau_{SS})b_j(\varepsilon) + \Upsilon \\ &\qquad c > 0, \qquad k' \geq k_b. \end{aligned} \tag{7}$$

The problem for a working-age household is

$$\begin{aligned} V_j^W(k, \varepsilon) &= \max_{c,h,k'} u(c, h) + \beta \left[(1 - \psi_a) \sum_{\varepsilon' \in \mathcal{E}} \pi_j(\varepsilon, \varepsilon') V_j^W(k', \varepsilon') + \psi_a V_j^R(k', \varepsilon) \right] \\ &\text{s.t.} \\ &(1 + \tau_c)c + k' = (1 + (1 - \tau_k)r)k + y_j(h, \varepsilon) - T_h[y_j(h, \varepsilon)] - \min\left[\tau_{SS} \cdot y_j(h, \varepsilon), \bar{t}_{SS}\right] + \Upsilon \\ &\qquad c > 0, \qquad k' \geq k_b, \qquad h \in [0, 1). \end{aligned} \tag{8}$$

The solutions of the dynamic programs (7) and (8) yield the decision rules for household

choices for consumption, savings, and labor supply, $\left\{g^a_{j,c}(k,\varepsilon), g^a_{j,k}(k,\varepsilon), g^a_{j,h}(k,\varepsilon)\right\}_{\substack{a\in\{W,R\}\\ j\in\{u,s\}}}$.

2.2 Definition of Equilibrium

Agents are heterogeneous at each point in time in the state $x \in X$. The agents' distribution over the states x is described by a measure of probability Γ_t defined on subsets of the state space X. Let $(X, \mathcal{B}(X), \Gamma_t)$ be a probability space, where $\mathcal{B}(X)$ is the Borel σ-algebra on X. For each $\omega \subset \mathcal{B}(X)$, $\Gamma_t(\omega)$ denotes the fraction of agents who are in probability state ω. There is a transition function $M_t(x, \omega)$ that governs the movement over the state space from time t to time $t+1$ and that depends on the invariant probability distribution of the idiosyncratic shock $\bar{\pi}_j(\varepsilon)$ and on the decision rules obtained from the household problem.

Definition 1 (Competitive economic equilibrium)**.** *Given initial conditions K_1 and Γ_1, a competitive economic equilibrium is a sequence of Social Security tax rates and benefit schedules $\{\tau_{SS}, b_j(\varepsilon_t)_{j\in J}\}_{t=1}^{\infty}$, government expenditures and debt $\{G_t, B_t\}_{t=1}^{\infty}$, lump-sum transfers $\{\Upsilon_t\}_{t=1}^{\infty}$, tax policies $\{\boldsymbol{\tau}_t\}_{t=1}^{\infty}$, value functions $\{V_t(x), g_{c,t}(x), g_{k,t}(x), g_{h,t}(x)\}_{t=1}^{\infty}$, factor prices $\{r_t, w_t\}_{t=1}^{\infty}$, firm plans $\{K_t, N_t\}$, average earnings $\{AE_t\}_{t=1}^{\infty}$, and measures $\{\Gamma(x)\}_{t=1}^{\infty}$ such that, $\forall t$*

1. *Given factor prices, taxes, and transfers, $\{V_t(x), g_{c,t}(x), g_{k,t}(x), g_{h,t}(x)\}$ solve the household problems in (7) and (8).*

2. *Given factor prices, $\{K_t, N_t\}$ satisfy equations (2) and (3).*

3. *Markets clear:*

   (a)
   $$A_{t+1} = \int g_{k,t} d\Gamma_t(x) = K_{t+1} + B_{t+1}$$

   (b)
   $$Y_t = \int g_{c,t}(x)\, d\Gamma_t(x) + K_{t+1} - (1-\delta) K_t + G_t$$

   (c)
   $$N_t = \sum_{j\in\mathcal{J}} \int exp(\varepsilon_t)\, g_{j,h,t}(k,\varepsilon)\, d\Gamma^W_{j,t}(k,\varepsilon)$$

4. *The government budget constraint clears*

$$G_t + \Upsilon_t + (1+r_t)B_t = \sum_{j\in\mathcal{J}} \int T_t\left(y^W_{j,t}(k,\varepsilon)\right) d\Gamma^W_{j,t}(k,\varepsilon) + \tau_{k,t} r_t \int_X k d\Gamma_t(x) + \tau_{c,t}\int_X g_{c,t}(x)\, d\Gamma_t(x) + B_{t+1}$$

5. *The Social Security budget balances*

$$\sum_{j\in\mathcal{J}} \int b_j(\varepsilon)\, d\Gamma^R_{j,t}(k,\varepsilon) = \int \min\left[\tau_{SS} y^W_{j,t}(k,\varepsilon), \bar{t}_{SS}\right] d\Gamma^W_{j,t}(k,\varepsilon).$$

6. *We can split Γ_t into the invariant distributions, $\Gamma^W_j(k,\varepsilon)$ and $\Gamma^R_j(k,\varepsilon)$. For any $\omega \in \mathcal{B}(\mathcal{K}\times\mathcal{E})$, distributions $\Gamma^W_j(k,\varepsilon)$ and $\Gamma^R_j(k,\varepsilon)$ are consistent with household decisions. Meaning that for all $j \in J$,*

$$\begin{aligned}
\Gamma^W_{j,t}(\mathcal{K},\mathcal{E}) =& (1-\psi_a)\int \sum_{\varepsilon'\in\mathcal{E}} 1_{\left\{g^W_{j,k}(k,\varepsilon)\in\mathcal{K}\right\}} \pi_j(\varepsilon,\varepsilon')\, d\Gamma^W_j(k,\varepsilon) \\
&+ \psi_d \int \sum_{\varepsilon\in\mathcal{E}} \overline{\pi}_j(\varepsilon) 1_{\left\{g^R_{j,k}(k,\varepsilon)\in\mathcal{K}\right\}} d\Gamma^R_{j,t}(k,\varepsilon) \\
\Gamma^R_{j,t}(\mathcal{K},\mathcal{E}) =& (1-\psi_d)\int 1_{\{\varepsilon\in\mathcal{E}\}} 1_{\left\{g^R_{j,k}(k,\varepsilon)\in\mathcal{K}\right\}} d\Gamma^R_{j,t}(k,\varepsilon) \\
&+ \psi_a \int 1_{\{\varepsilon\in\mathcal{E}\}} 1_{\left\{g^W_{j,k}(k,\varepsilon)\in\mathcal{K}\right\}} d\Gamma^W_{j,t}(k,\varepsilon)
\end{aligned}$$

where the conditional transitions $M^a_{j,t} : (\mathcal{K}\times\mathcal{E}, \mathcal{B}(\mathcal{K}\times\mathcal{E})) \to (\mathcal{K}\times\mathcal{E}, \mathcal{B}(\mathcal{K}\times\mathcal{E}))$ are explicitly written inside the sums.

3 Calibration

Demographics. In the model, households expect to work for $J_W = 40$ years and to stay retired for $J_R = 15$ years. The aging and death probabilities are hence the inverse of those years, $\{1/J_W, 1/J_R\}$, respectively. The fraction of the population with high skill, μ_s, is set to 41 percent, to match the share of the population with a college degree in the US from Kindermann and Krueger (2022).

Preferences. The period utility is

$$u(c,h) = \frac{c^{1-\gamma}}{1-\gamma} - \theta \frac{h^{1+\frac{1}{\varphi}}}{1+\frac{1}{\varphi}} \tag{9}$$

where γ is relative risk aversion, θ controls the preference of labor vs. consumption, and φ is the inverse of the Frisch elasticity. We set $\gamma = 2.0$ and $\varphi = 2.0$. The value of θ is obtained by calibrating it to target average working hours in the model of 30 percent of disposable time. The discount factor β is calibrated to match a capital-output ratio of 3.0.

Technology. We set $\alpha = 36$ percent to match capital's share of national income. We calibrate Z so that GDP, Y, is normalized to one.[6] We calibrate the depreciation rate of capital δ to be 5 percent so that investment is 15 percent of GDP.

Labor Income. We set the college skill premium, ζ_s/ζ_u, to 175 percent to match the data for the US. For each skill, we discretize the productivity values into seven points, in increasing order $\varepsilon_1 < ... < \varepsilon_7$. In each group, the first five productivity levels are drawn from distinct skill-dependent Markov processes, while the highest two are independent of education. For the first "regular" worker states, productivity follows an AR(1) process in logs:

$$\log \varepsilon' = \rho_j \log \varepsilon + \iota,\ \iota \sim N\left(0, \sigma^2_{\varepsilon,j}\right). \tag{10}$$

We discretize this process into a Markov chain with realizations $\{\varepsilon_1, ..., \varepsilon_5\}$ and transition matrix $[\pi_{i,j}]_{i,j=1,5}$.

We follow Carroll and Hur (2023) and use their estimated values for the persistence ρ_j and standard deviation $\sigma_{\varepsilon,j}$ obtained from the PSID. The values are $\rho_h = 0.941$ and $\sigma_{\varepsilon,u} = 0.197$ and $\rho_s = 0.914$ and $\sigma_{\varepsilon,s} = 0.229$.

We name the sixth and seventh productivity elements the "steppingstar" and "superstar" states, respectively. At those states workers have a substantially higher productivity than the average worker, with the "super-star" state being an extreme outlier. Following Kindermann and Krueger (2022), we calibrate the associated parameters of the discretized Markov transition matrix to match moments at the top of the earnings and wealth distributions. In particular, we calibrate the set of transition probabilities, $\{\pi_{x,6}, \pi_{6,6}, \pi_{6,7}, \pi_{7,7}\}$, and the

[6] This normalization is convenient, as it allows us to reduce the number of variables we need to calibrate.

associated productivity shock levels, $\{\varepsilon_6, \varepsilon_7\}$, to match the top 5 percent and top 1 percent shares and Gini coefficients of the earnings and wealth distributions. For those moments, we target the values computed from the 2019 SCF update of Kuhn and Ríos-Rull (2015).

Government. We calibrate the benchmark consumption tax, τ_c, to 6.4 percent and the capital income tax τ_k is set to 27.3 percent, both being the average in the US between 1990-2000 as calculated by Carey and Rabesona (2003) using OECD data. For the labor earnings tax, the parameter τ_y is calibrated to target an average labor tax rate of 21 percent as in McDaniel (2007) and the curvature parameter ν_y is calibrated to generate a top marginal tax rate of 37.9 percent, which is the average top marginal tax rate from 2002 to 2023 (Tax Policy Center, 2023).

We calibrate the share of the total amount of transfers to GDP, Υ/Y, to be 2.2 percent, which lies within the range of 1.3 to 2.7 percent identified in the data as reported and calculated in Luduvice (2024) from the CBO (CBO, 2019) and the White House's Office of Management and Budget (OMB, 2023). It is also close to the value of 2.3 percent reported in Guner et al. (2023a). We set the aggregate level of government spending, G, relative to GDP as 18 percent following Trabandt and Uhlig (2011). Finally, the stock of debt, B, is calibrated in the benchmark economy to clear the government budget constraint, leading to a share of GDP of 64.1 percent, close to the value of 63 percent from Trabandt and Uhlig (2011).

Social Security. The aggregate average labor earnings, AE, used in the benefits payment schedule is obtained endogenously in the solution of the model. The contribution cap $\bar{t}_{SS}$ that closes the budget of the Social Security system is equal to 45.4 percent. The replacement rates r_1, r_2, r_3 and bend points b_1, b_2 are calibrated from Social Security data as in Huggett and Parra (2010) and Kindermann and Krueger (2022).

3.1 Summary of Calibration

We summarize the information associated with the calibrated parameters in the sequence of tables below. In Table 1, one can find the exogenously calibrated parameters and their sources. Table 2 shows the endogenously calibrated parameters, the targeted moments associated with each of them, the source of those moments for their data counterparts, and the

value of those statistics computed for the model economy.

Table 1: Exogenously calibrated parameters.

	Parameter	Value	Target / Source
Demographics			
Working and retirement years	J_W, J_R	$\{40, 15\}$	Standard
Aging and death probabilities	ψ_a, ψ_d	$\{1/J_W, 1/J_R\}$	Standard
Fraction of pop. with college	μ_s	41%	Kindermann and Krueger (2022)
Preferences			
Relative risk aversion	γ	2.00	Standard
Inverse Frisch elasticity	φ	2.00	Standard
Technology			
Capital share	α	0.36	Standard
K depreciation rate	δ	0.05	Standard
Labor Income			
AR(1) non-college	$\{\rho_u, \sigma_{\varepsilon,u}\}$	0.941, 0.197	PSID (Carroll and Hur, 2023)
AR(1) college	$\{\rho_s, \sigma_{\varepsilon,s}\}$	0.914, 0.229	PSID (Carroll and Hur, 2023)
College skill premium	$\{\zeta_u, \zeta_s\}$	1.00, 1.75	Carroll and Hur (2023)
Government			
Consumption tax	τ_c	6.4%	Carey and Rabesona (2003)
Capital income tax	τ_k	27.3%	Carey and Rabesona (2003)
Payroll tax	τ_{SS}	12.4%	IRS
Government spending	G/Y	18%	Trabandt and Uhlig (2011)
Lump-sum transfer	Υ/Y	2.2%	CBO (2019); OMB (2023)
Social Security			
Replacement rates	$\{r_1, r_2, r_3\}$	$\{0.90, 0.32, 0.15\}$	Soc. Sec. data (Huggett and Parra, 2010)
Bend points	$\{b_1, b_2, b_3\}$	$\{0.21, 1.29, 2.42\}$	Soc. Sec. data (Huggett and Parra, 2010)

Notes: The table shows model parameters, their numerical values, targeted moments in the model economy, and their data sources.

Table 2: Endogenously calibrated parameters.

	Parameter	Value	Target	Data	Model
Preferences					
Discount factor	β	0.934	K/Y	3.0	3.0
Labor disutility	θ	62.032	Average hours	0.3	0.3
Technology					
Aggregate productivity	Z	0.747	Normalize GDP	-	1.0
Labor Income					
Avg. Labor Earnings	AE	0.880	-	-	0.880
Government					
Scale parameter of labor tax	τ_y	0.224	Avg labor tax rate	21%	21%
Curvature of income taxes	ν_y	0.132	Top mg. tax rate	37.9%	37.9%
Government Debt	B/Y	0.641	Balance govt budget	63%	64.1%
Social Security					
Contribution cap	$\bar{t}_{SS}$	0.450	Balance Soc. Sec. budget	-	-
Inequality Statistics					
Prob. of staying stepping-star	$\pi_{6,6}$	0.9698	Earnings 95% - 99%	18.4	17.9
Prob. to superstar	$\pi_{6,7}$	0.0009	Earnings 99% - 100%	18.8	20.2
Prob. to star region	$\pi_{x,6}$	0.0056	Earnings Gini	0.67	0.65
Stepping-star shock	ε_6	17.2212	Wealth 95% - 99%	27.4	24.2
Superstar shock	ε_7	1090.7770	Wealth 99% - 100%	35.5	27.0
Prob of staying superstar	$\pi_{7,7}$	0.9270	Wealth Gini	0.85	0.85

Notes: The table shows model parameters, their numerical values, and targeted moments in the model economy. The details for the data counterparts of the targets are outlined in the text. For the inequality statistics, the data moments are taken from the 2019 update of the SCF calculations by Kuhn and Ríos-Rull (2015).

4 Numerical Experiment

We now turn to the central numerical experiment of the paper. Our purpose is to uncover the indirect preferences for taxes and transfers over a wide menu of fiscal policies. Specifically, we compare policies with different consumption and capital income tax rates combined with different labor income tax schedules, the latter in terms of both average levels and degrees of progressivity, and identify the policy that maximizes average welfare.

We start the economy in the calibrated steady state and evaluate permanent fiscal policy reforms, assuming that the government can perpetually and fully commit to any policy that it enacts. As part of this process, we solve for the transition to the associated final steady state for each policy. Attention to the transition is necessary because some policies lead

to higher capital stocks in the long run relative to the initial capital level. Building to these higher capital levels requires an extended period of increased investment and reduced consumption. Households that must undergo this stage of the transition may find that the short-run costs from forgone consumption outweigh the long-run benefits from having more capital. Of course, the opposite scenario is equally important: the welfare gains from eating into an initially high capital stock are realized early in the transition while the costs from low capital come later. Given the life-cycle properties of our model, restricting the analysis to steady-state comparisons would be misleading, because the costs will be borne by those currently alive, while the gains would accrue to the yet-to-be born.[7]

Let $\mathcal{P}$ denote the menu of fiscal policies, p. p has five elements: four permanent tax parameters and one time-varying transfer, that is,

$$p \equiv \{(\tau_y, \nu_y, \tau_k, \tau_c), \Upsilon_t\}_{t=1}^{\infty}. \tag{11}$$

The time-varying path of transfers results from the government balancing its budget in each period of the transition.[8]

Let $V_x(\overline{p})$ be the indirect utility from tax reform $\overline{p}$ for a household with initial state vector $x \equiv \{k, \varepsilon, j, a\}$, and $p_x^\star$ be the household's *most preferred policy*, given by

$$p_x^\star = \arg\max_{\overline{p} \in \mathcal{P}} V_x(\overline{p}). \tag{12}$$

Define p_{SP} as the policy that maximizes social welfare measured by the population-weighted sum of the indirect utilities of all households that are alive in the initial steady state. Formally,

$$p_{SP} = \arg\max_{\overline{p} \in \mathcal{P}} \int V_x(\overline{p}) d\Gamma_0(x), \tag{13}$$

where Γ_0 is the initial wealth distribution.

In our computational experiment, we allow a wide range of rates for each of the tax instruments. Specifically, we construct the menu of fiscal policies, $\mathcal{P}$, from combinations of four defined grids over the different tax rates. We assign 6 potential values for the rates

[7] This tension can be seen in Conesa and Garriga (2003), who study why Social Security reforms that raise long-run welfare are not implemented because they do not benefit the current generations.

[8] The level of debt is fixed exogenously in our baseline. Later, when the debt-to-GDP ratio is held constant instead, the level of debt will also be time varying.

$\{\tau_y, \tau_k, \tau_c\}$ and 18 for the curvature parameter that governs the progressivity level, ν_y. The grids are as follows:

$$\begin{gathered}
\tau_y \in \{2.0\%, 12.0\%, 22.4\%, 32.0\%, 42.0\%, 57.0\%\}, \\
\nu_y \in \{0.0\%, 2.7\%, 6.5\%, 10.2\%, 11.7\%, 13.2\%, 14.7\%, 17.0\%, 20.2\%, \\
22.2\%, 23.2\%, 24.2\%, 25.2\%, 26.2\%, 27.2\%, 28.2\%, 29.2\%, 30.2\%\}, \\
\tau_k \in \{0.0\%, 10.0\%, 18.7\%, 27.3\%, 40.0\%, 60.0\%\}, \\
\tau_c \in \{0.0\%, 3.2\%, 6.4\%, 12.8\%, 25.6\%, 51.2\%\}.
\end{gathered}$$

Each tax grid is centered around the corresponding calibrated value for the benchmark economy shown in Table 1, and, together, the grids span a wide range of tax policies in terms of both tax levels and composition. For instance, eliminating capital income taxation and consumption taxation is possible ($\tau_k = 0$ and $\tau_c = 0$, respectively), and although labor income taxation cannot be eliminated entirely since τ_y is strictly positive, it is possible to come very close.[9] Second, labor income taxation can range from linear ($\nu_y = 0$) with a flat rate of τ_y on all households to something approximating the two-tiered policy in Conesa et al. (2009), where households below an income threshold are exempted while those above face a high, flat marginal rate, with the top value more than double the calibrated value ($\nu_y = 30.2\%$). Because the maximum values along all four grids are quite high, a variety of stringent taxation scenarios is permitted as well. Due to its importance for redistribution and the economy's allocation sensitivity to it, the grid for progressivity contains more points than the other grids. We have purposefully made the progressivity grid finer by adding points in the neighborhood of the optimal value.

We exclude negative transfers from consideration, both in the terminal steady state and in any period along the transition. Cases that are found to produce $\Upsilon_t < 0$ are deemed to be inadmissible and discarded. We justify this restriction in three ways. First, from an empirical standpoint, "head taxes" are not relevant. Second, in a model with binding borrowing limits, negative lump-sum taxes could require that some households default on their tax obligations in order to maintain positive consumption. Since we do not wish to model default, we simply ignore these cases. Finally, previewing our results, all policies with

[9]The model does not require $\tau_y > 0$. Rather, the decision to restrict τ_y in this manner was made after initial results showed that eliminating labor income taxation almost always produced negative transfers. Imposing a minimum of 2 percent produced a considerably greater number of feasible policies while still capturing the spirit of very low labor income taxation.

near-zero lump-sum transfers – whether measured in present value or in long-run levels – belong to the group of policies with the lowest welfare levels. From the 3888 tax policy combinations in $\mathcal{P}$, we find 2543 permissible options. Given that we find strong support for high positive transfers, this restriction ultimately does not matter for our results.

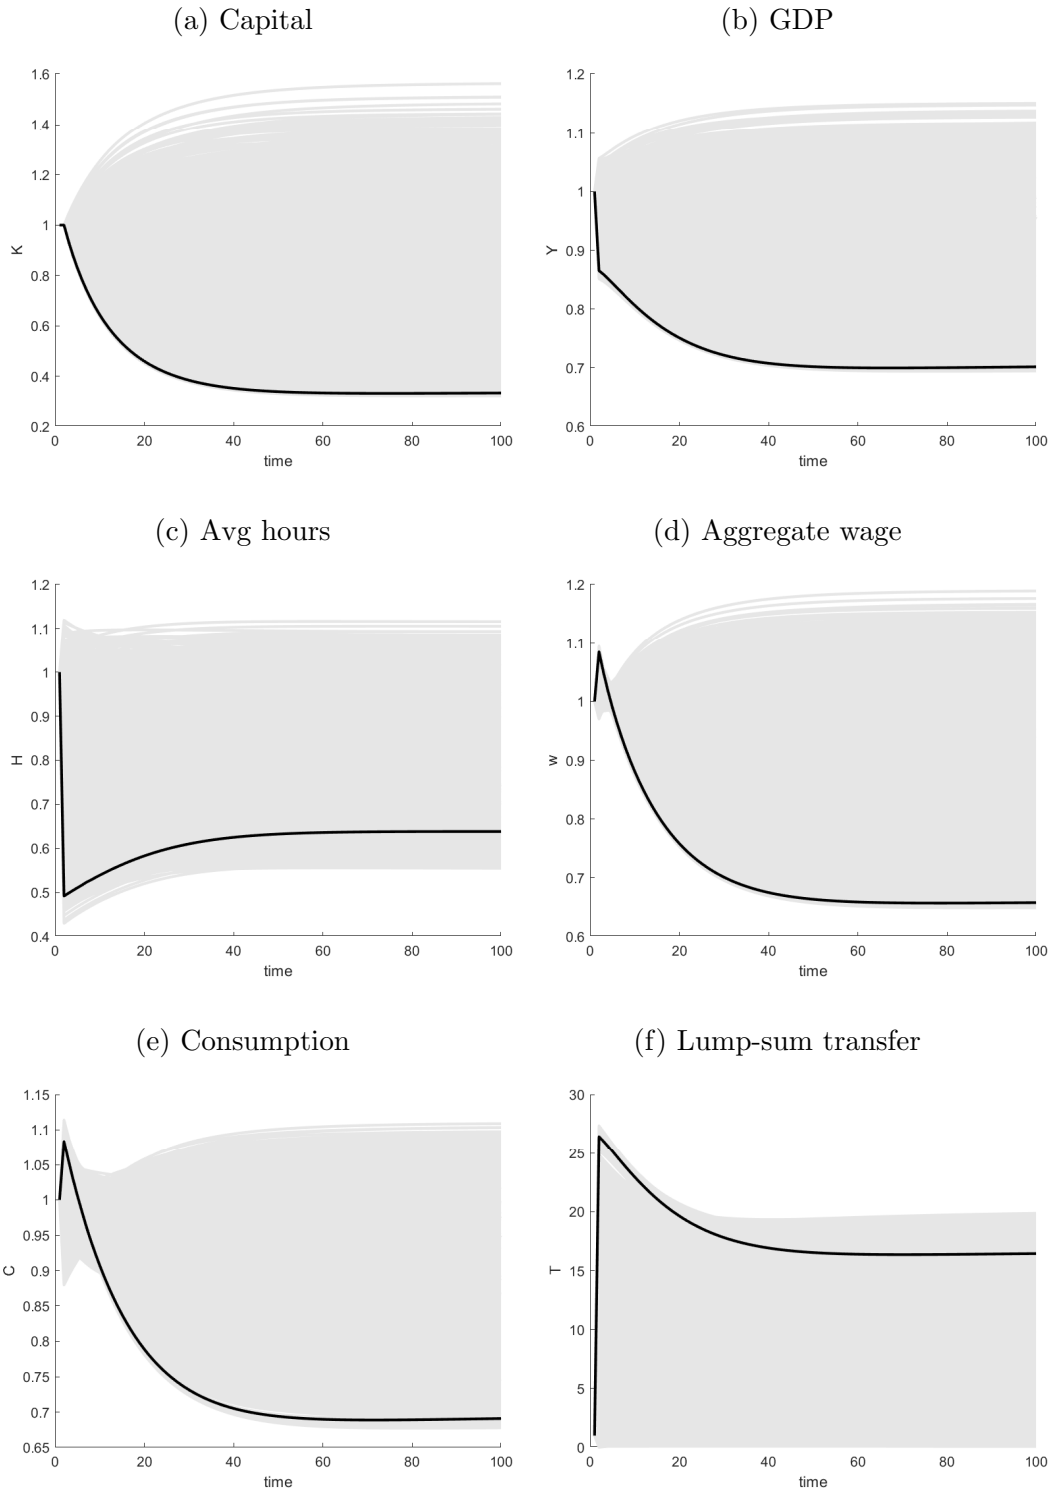
4.1 Aggregate Effects of Tax Changes

The policies in $\mathcal{P}$ span a wide variety of equilibrium transitional dynamics. The fan chart in Figure 1 plots all the feasible transition paths for the capital stock, GDP, hours, the aggregate wage, aggregate consumption, and transfers relative to their pre-reform levels. The darkened line in each figure highlights the path associated with $p_{SP} = \{57.0\%, 22.2\%, 60.0\%, 51.2\%\}$.[10] The consequences of following the optimal policy are striking: p_{SP} tanks the economy. GDP drops by nearly 30 percent over the transition, with about half of that decline occurring in the first period (panel b). This steep decline is driven primarily by a large fall in hours worked (panel c) early in the transition, the result of the combination of high distortionary taxation on labor and the wealth effect from large lump-sum transfers. Aggregate hours recovers partially in the long run, as high-income households respond to a larger negative wealth effect by working more. Very high capital income taxation contributes to the decline in GDP through severe capital shallowing (panel a).[11] In the long run, the capital stock is only one-third of its original level. Viewed against the collection of alternatives, p_{SP} is among a handful of high-tax policies that produce the lowest aggregate activity.

[10] Notice that three of the elements are at the upper limit on their grid. Expanding these grids introduces substantial computational difficulties, reducing the number of computationally and economically feasible cases. We have implemented a limited number of extra robustness cases, and we discuss the results in Appendix C.

[11] Investment is negative in the early periods of transition. Imposing capital irreversiblity to avoid this, however, would only smooth out the decline in aggregate activity, not prevent it. A non-negativity requirement on aggregate investment makes the capital stock less elastic and therefore an even more attractive target for taxation.

Figure 1: Transition paths of aggregates.



Notes: The figure shows the transitional dynamics for aggregate capital, output, average hours, aggregate wage, aggregate consumption, and aggregate transfers for all the feasible tax menus. In all panels, the solid black line shows the path induced by p_{SP}. The initial period represents the original steady-state quantities, which are normalized to 1.0. The duration of the transition is truncated at 100 years for the sake of exposition.

Despite the extreme negative effect that p_{SP} has on aggregate activity, because of the high degree of inequality in the initial steady state, a majority of households support the reform. Most households are below average in income, and wealth is highly concentrated in the right tail. As a result, while the reform greatly reduces the size of the economic pie overall, most households end up with a larger absolute slice. This is apparent in the rise in average consumption in the early transition (panel (e)). Since this front-loaded boost in consumption is combined with increased leisure and a much larger transfer, it is not surprising that the welfare effects are large and positive on average.

A strong desire for transfers is a common theme in incomplete markets models that are calibrated to the US wealth distribution. In recent papers that allow for more flexibility in the menu of taxes chosen by the planner, such as in Dyrda and Pedroni (2023), the optimal time-dependent transfer is 40 percent of output in the initial period of the transition. A similar pattern is present in Ferriere et al. (2022), where the optimal level of transfers is 45 percent of median income, driven mostly by demand in the left tail of the wealth distribution. In both cases, these large transfers are supported with high rates of distortionary taxation, in ranges similar to ours. In another variation of the same finding, Boar and Midrigan (2022) show that with an optimal high and flat labor income tax, lump-sum transfers play a significant role in achieving utilitarian welfare gains and have a better redistribution effect than increasing marginal income taxes.

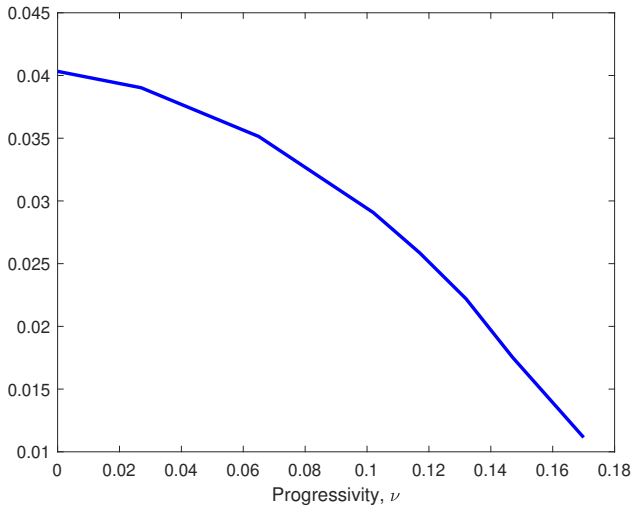
This paper generalizes these high-transfer high-tax fiscal policy prescriptions by allowing for more tax instruments and for more dimensions of household heterogeneity. Given sufficient flexibility for raising tax revenue, the motivation for redistribution dominates efficiency considerations in the workhorse macroeconomic model of inequality. In particular, the desire for insurance is so strong that it is optimal to elevate tax rates in order to finance a large lump-sum transfer. This alleviates the welfare cost of hitting the borrowing constraint and takes advantage of the high degree of inequality in the initial wealth distribution. In this regard, permitting the government to tax multiple bases through capital, consumption, and labor income taxes is essential to reveal the intensity of the underlying incentive to redistribute.

Then, from within this broader agenda, progressive labor income taxation further shifts the overall tax burden away from households with high marginal utility of consumption and partially insures households against the fundamental source of risk in the model: persistent

idiosyncratic labor income shocks.

Progressivity and Revenues. How government revenue responds to an increase in progressivity is not immediately obvious. On the one hand, revenue could rise since more progressivity puts higher tax rates on households with higher labor income, which would increase revenue (all other things being equal). On the other hand, it also shifts the incentive to work additional hours from the most productive toward the least, meaning that aggregate labor input will decline, and thus revenue may fall. In this environment, the second effect dominates: more progressive labor income tax schedules are associated with lower revenues and lower transfers. Keeping other tax rates fixed, the present discounted value of lump-sum transfers decreases as progressivity, ν_y, increases (Figure 2).

Figure 2: Transfer as a function of tax progressivity.



Notes: The figure shows the present discounted value of equilibrium transfers as a function of the progressivity parameter, ν_y. All other tax parameters are fixed at their initial steady-state values. Without loss of generality, the PDV shown is calculated using the formula for the initial young as described in Appendix A.

Effects on Inequality. The menu of tax policies produces a wide range of possible inequality effects. Table 3 reports the lower and upper bound on the Gini indices of the wealth, earnings, and consumption distributions, along with the value for each at the optimal tax menu, p_{SP}. While the earnings and wealth Ginis are in line with the data, the model's Gini coefficient of consumption is higher than commonly estimated from the CEX and PSID data (see Krueger and Perri, 2006 and Krueger et al., 2016). We do not view this discrepancy as

a problem. First, there is reason to think that the true Gini coefficient is likely higher than in those datasets. As noted in Heathcote et al. (2010, 2023), these datasets under-represent the top tails of the income, wealth, and consumption distributions. Moreover, high-income households have larger shares of expenditures in luxury goods which further exacerbates consumption inequality (Aguiar and Bils, 2015). Second, because our model features long right tails both in earnings and especially in wealth, without non-homotheticities to dampen consumption it will produce large consumption Gini coefficients (Ferraro and Valaitis, 2024). However, if "steppingstar" and "superstar" households are excluded from the calculation, as they likely are from the data, the consumption Gini is only 0.32, which is consistent with the magnitude estimated using the CEX and PSID data.

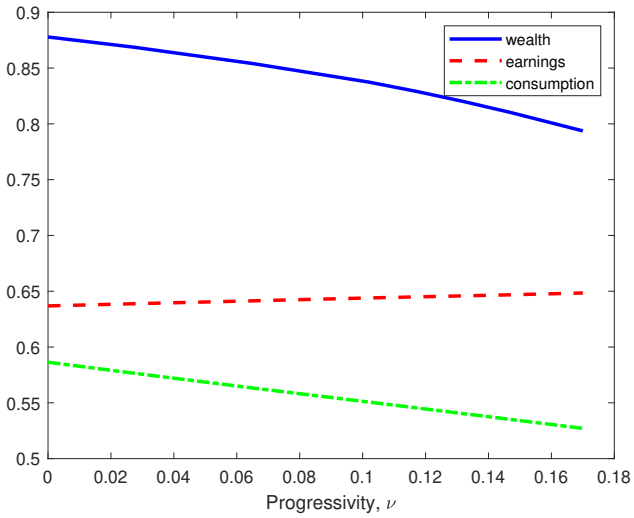
Table 3: Range of Gini indices.

Gini	Initial SS	Minimum	Maximum	p_{SP}
Wealth	85	68.8	95.9	88.5
Earnings	65	62.3	77.8	76.4
Consumption	54	28.3	60.3	30.2

Notes: The table shows the minima and maxima Gini values across all feasible (i.e., non-negative transfers) tax policies, along with the those associated with the optimal tax policy.

Because the tax menu has many points in the progressivity dimension, the marginal effect of changing this parameter can be readily seen. Figure 3 depicts how the Ginis for wealth, earnings, and consumption change with progressivity when the other three tax parameters remain fixed at their initially calibrated values. Across the feasible cases in this region of the tax parameter space, movements in ν_y have relatively larger impacts on consumption and wealth inequality than on earnings inequality. Pre-tax earnings inequality remains relatively flat across the range of feasible values for ν_y, while wealth and consumption inequality decrease as we increase progressivity.

Figure 3: Long-run Gini coefficient across progressivity.



Notes: The figure shows the long-run Gini coefficients of wealth, earnings, and consumption as a function of the progressivity parameter, ν_y. All other tax parameters are fixed at their initial steady-state values.

4.2 Household Preferences over Tax Policies

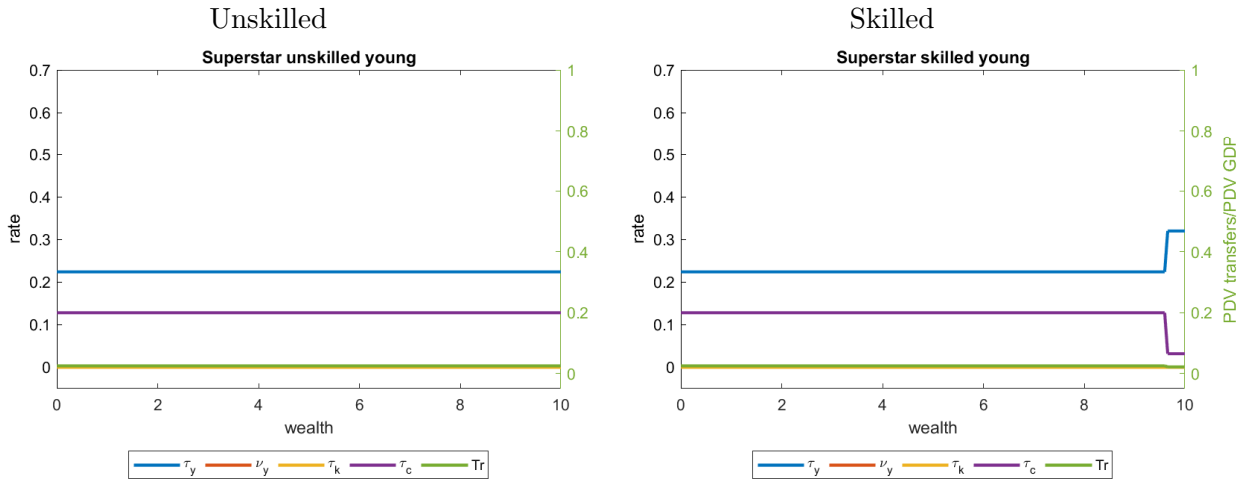
Next, we put a magnifying lens over the state space to examine how preferences over the available tax menus changes with household characteristics. The model features a significant amount of household heterogeneity. Not only are households very unequal in their levels of income and wealth, but they also differ in the composition of their income between labor and capital. In addition, because retirement benefits are fixed, old households face no risk around their future income stream. These factors lead to considerable disagreement over what policy should be. Although the variation in policy preferences can appear complicated at first glance, each household's $p^\star$ comes down to a trade-off between maximizing the transfer while minimizing its own tax burden.

Effect of Wealth. Wealth heterogeneity is a central ingredient in the model and also the dimension along which the widest disagreement occurs. As a rule of thumb, low-wealth households want a high transfer. To achieve this transfer, they prefer to impose high tax rates on consumption and capital income. The motivation for high capital income taxes is simple: these households own a small share of the capital stock and most have a low probability of being wealthy in the near future. The motivation for high consumption taxation is related.

The jump in consumption taxes to a new permanent level acts as a tax on initial wealth, which is highly unequally distributed; given that the tax is constant after the initial period, there is no further intertemporal aspect to consumption taxes.

This preference does not imply that low-wealth households face no drawbacks from these high-tax policies. First, flat consumption taxation is more onerous for poor households, especially those near the borrowing limit, since they have higher marginal propensities to consume. Second, high capital income taxation discourages the accumulation of capital in the future, and this decline, in general equilibrium, puts downward pressure on wages. For all but a few near-average-wealth households, however, these considerations are of second-order importance relative to the higher transfer. Many low-wealth households also support policies with a high average labor tax rate; however, the strength of support depends on other factors like productivity (or, equivalently, initial and expected wages), skill, and the progressivity of the tax schedule. We discuss the role of these factors later. For now, holding fixed other factors, as the initial wealth level of a household increases, the desire for a high transfer wanes and support increases for reducing taxes on consumption and on capital income.

Figure 4: Most preferred policy of young superstars.



Notes: The figure shows the most preferred level of each of the tax instruments along the wealth dimension for young superstar workers. The left panel highlights unskilled workers and the right panel skilled workers. For both panels, the y-axis on the left-hand side marks the different tax rates, while the y-axis on the right-hand side shows the present discounted value of the lump-sum transfers as a fraction of the present discounted value of output. For more details on how to compute the discounting of aggregate measures, see Appendix A. All lines are smoothed for exposition purposes using a moving median adjustment.

At high wealth levels, households prefer low transfers that are financed entirely through

labor income taxes. Due to the wealth effect on leisure, labor income as a fraction of total income declines as households get wealthier. By financing the transfer only with labor income taxes, these households avoid taxation while still securing a bit of extra consumption for themselves. This logic is most apparent in the plot of young superstar households' ideal tax structure, which is shown in Figure 4. All superstars, regardless of their wealth or skill, want to eliminate transfers and raise only enough revenue to cover government spending and debt. Within either skill group, preferences for how to raise this revenue depend only on differences in initial wealth. Those with sufficiently high wealth would like to tax labor income only. Superstars who do not meet this condition still expect to be wealthy very soon, so they want low capital income and consumption taxes; however, they have a very high wage and a strong desire to work in the short run, so they do not want to place the entire tax burden on labor income.

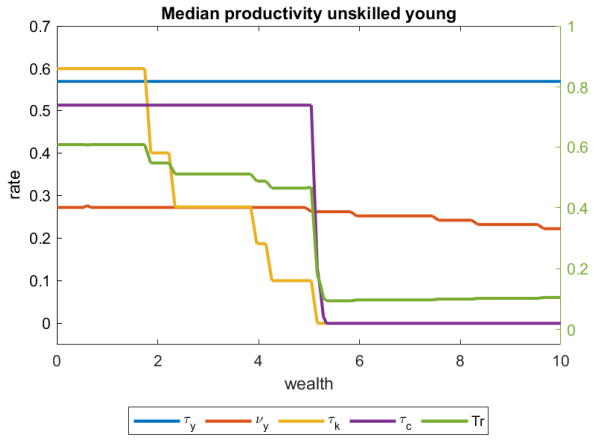
Effect of Skill Type. Skill is the most straightforward dimension in terms of its effect on policy preferences. Skilled households behave for all intents and purposes like a higher-wage unskilled household.[12] Comparing panels (a) and (b) in Figure 5, we see that, across the wealth distribution, $p^\star$ is broadly similar for the unskilled and the skilled. At low wealth levels, both types support a mixture of taxes and high transfers. The switch to labor-tax-only government financing occurs at a slightly lower level of wealth for the skilled, but again the pattern across wealth is very similar. The main disagreement between the two groups is over progressivity. The unskilled generally favor some moderation of progressivity at higher wealth levels as it mildly increases the transfer level. Skilled households, which all else equal have higher earnings, would trade away some of the transfer to maintain high levels of progressivity.

Effect of Age. A key dimension of heterogeneity in our model is that households do not expect to live forever. The age dimension affects households' most preferred tax menu and the resulting lump-sum transfer that it would finance as shown in Figure 6. If we compare with Figure 5, we can observe that workers, when compared to their retired counterparts, are willing to trade higher consumption taxes and progressivity for lower capital income taxes along the wealth distribution. Retirees favor modest to zero progressivity, i.e., a flat labor

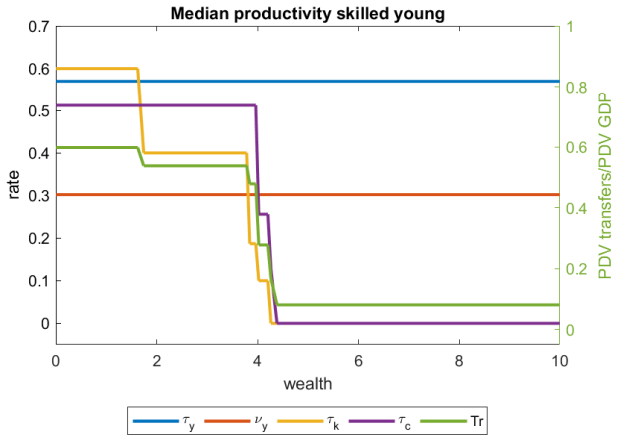
[12]Beyond just productivity levels there are other small differences arising from the slightly higher variance in the stochastic process for skilled labor productivity.

Figure 5: Most preferred policy by skill.

(a) Unskilled young



(b) Skilled young



Notes: The figure shows the most preferred level of each of the tax instruments along the wealth dimension for workers with regular productivity levels. The left panel highlights unskilled workers and the right panel skilled workers. For both panels, the y-axis on the left-hand side marks the different tax rates, while the y-axis on the right-hand side shows the present discounted value of the lump-sum transfers as a fraction of the present discounted value of output. For more details on how to compute the discounting of aggregate measures, see Appendix A. All lines are smoothed for exposition purposes using a moving median adjustment.

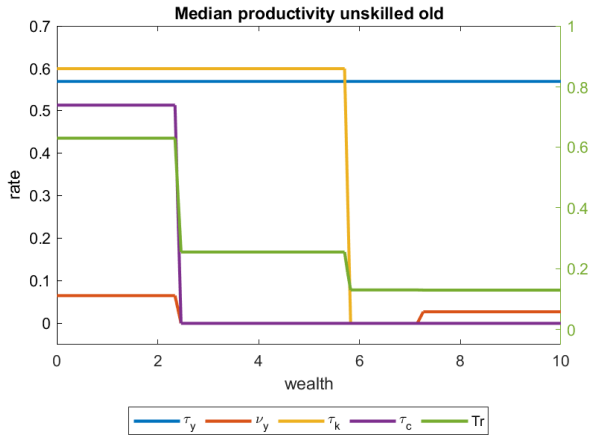
tax on workers, and all but the richest retirees favor using high capital income taxes in order to raise revenue for redistribution.

Effect of Initial Wage/Productivity. A large portion of the disagreement over policy comes from differences in the initial distribution of wages (i.e., labor productivity) over young households. Excluding stepping-stars and superstars, young worker households comprise about two-thirds of the model economy. Figure 7 plots the most preferred policy of two types of these households: one with the lowest wage and one with the highest wage (again, excluding star households). The strongest disagreement along this dimension is over how labor income is taxed, and consequently, the level of transfers.

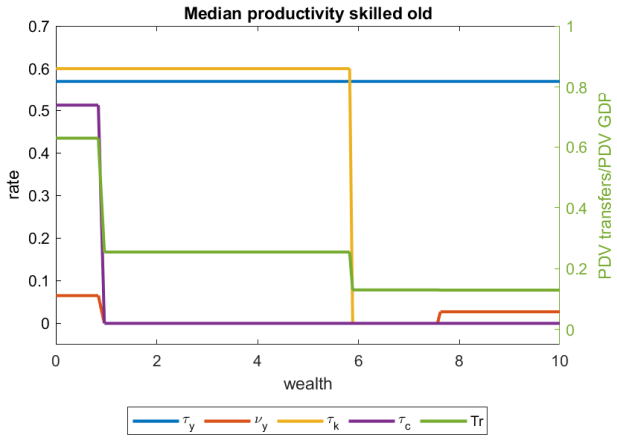
Starting with the low-wage earner (panel 7a), at low levels of initial wealth, the transfer is the most important factor. These households want the transfer to be as large as possible and are willing to face a linear labor tax to achieve it. At sufficiently high initial wealth levels, these households give back some transfer to "purchase" progressivity in the labor tax. Notice that even at very high levels of wealth, the low-wage households still favor just moderate progressivity. The wage process has substantial persistence, so that even initially

Figure 6: Most preferred policy of old.

(a) Unskilled old



(b) Skilled old



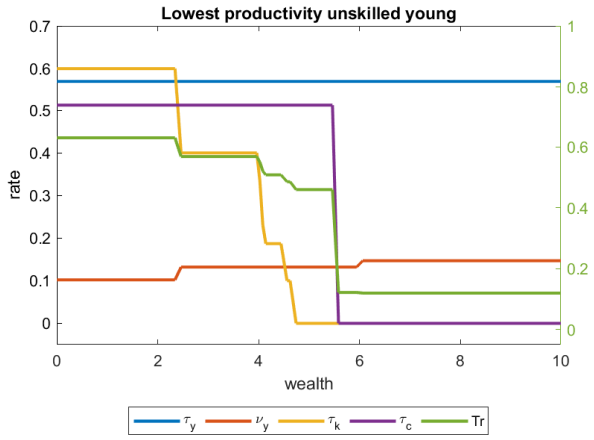
Notes: The figure shows the most preferred level of each of the tax instruments along the wealth dimension for retired workers who had regular productivity levels during their working years. The left panel highlights unskilled retirees and the right panel skilled retirees. For both panels, the y-axis on the left-hand side marks the different tax rates, while the y-axis on the right-hand side shows the present discounted value of the lump-sum transfers as a fraction of the present discounted value of output. For more details on how to compute the discounting of aggregate measures, see Appendix A. All lines are smoothed for exposition purposes using a moving median adjustment.

rich households expect to run down their savings well into the future if they are currently low productivity. Anticipating that they will eventually be poor, they wish to maintain a moderate social safety net.

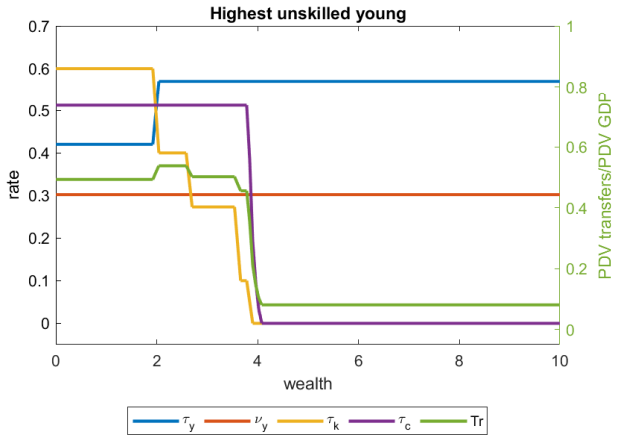
High initial wage earners, on the other hand, have a weaker indirect preference for transfers. Those with low initial wealth want lower average labor taxes and high progressivity, since they plan to exploit their good fortune by working a lot and accumulating wealth. Like their low-wage counterparts, wealthy high-wage households want to maintain some positive level of transfers, but not at the expense of reducing progressivity.

Figure 7: Most preferred policy by initial wage.

(a) Lowest wage (unskilled)



(b) Highest non-star wage (unskilled)



Notes: The figure shows the most preferred level of each of the tax instruments along the wealth dimension for unskilled workers. The left panel highlights workers with the lowest regular wage and the right panel the ones with the highest regular wage. For both panels, the y-axis on the left-hand side marks the different tax rates, while the y-axis on the right-hand side shows the present discounted value of the lump-sum transfers as a fraction of the present discounted value of output. For more details on how to compute the discounting of aggregate measures, see Appendix A. All lines are smoothed for exposition purposes using a moving median adjustment.

Distribution of Preferences. While the figures above relate the wide range of households' indirect preferences over fiscal policy by initial state, they are not quite informative about the initial distribution of households over those states. For instance, from Figure 7 panel (a), we know that a sufficiently wealthy, unskilled young household with low labor productivity does not want to tax consumption or capital income, but this preference has no bearing on equilibrium policy because there is zero mass in that region of the state space. To get a sense of how any particular type of household affects policy outcomes, it is helpful to know the policy this group favors along with the population share of the group. This information is displayed along with that from other groups and for different quantiles of the wealth distribution in Table 4.

After accounting for the distribution of these subgroups, the strong preference for redistribution becomes even more apparent. The tax instruments that generate disagreement are the average tax rate, τ_y, and the progressivity value, ν_y. In that case, the following facts hold: for retired households, the outcome is nearly a high flat labor income tax; for "star" households, the outcome is a moderate flat labor income tax; and for young non-star

households, there is a contest between progressivity and the average labor tax, along the lines of the aforementioned mechanism of using progressivity as a means of generating more revenue and hence larger lump-sum transfers. The taxes on consumption and capital are all at their highest levels. When we break down tax preference along the wealth distribution, we observe that the top 20 percent would prefer a higher level of progressivity than the benchmark paired with absence of taxes on both capital and consumption, making a choice geared toward the protection of their accumulated wealth. As we walk through the table and account for the population share of each group, the trade-off faced by the planner when weighing the disagreement on the progressivity rate between groups is apparent, making it the only "interior" choice of all tax instruments.

Table 4: Social welfare maximizing policy by subgroup.

HH Type	τ_y	ν_y	τ_k	τ_c	Υ/Y	Population Share
Young, non-star						
unskilled	57.0	28.2	60.0	51.2	60.4	36.9
skilled	57.0	30.2	60.0	51.2	59.8	25.7
All stars	32.0	0.0	10.0	0.0	2.1	14.0
Retired						
unskilled	57.0	6.5	60.0	51.2	62.8	13.8
skilled	57.0	6.5	60.0	51.2	62.8	9.6
Wealth						
Bottom 50%	57.0	17.0	60.0	51.2	62.6	50.0
Mid 50% − 80%	57.0	30.2	60.0	51.2	59.8	30.0
Top 20%	57.0	25.2	0.0	0.0	9.7	20.0

Notes: The table shows the tax rate that maximizes welfare for each of the taxes in the available menu across the different household types and for different quantiles of the wealth distribution. The last column shows the population share that the group of household type represents. All units are in percents. Υ/Y is the ratio of the present discounted value of the lump-sum transfers to GDP.

5 Optimality

Under the utilitarian specification, the social planner effectively places more weight on households with high marginal utility, so it is not surprising that p_{SP} aligns with the preferences of unskilled workers. The optimal policy calls for a massive redistribution of income. The present discounted value of transfers-to-GDP over the transition is 61.1 percent, up from

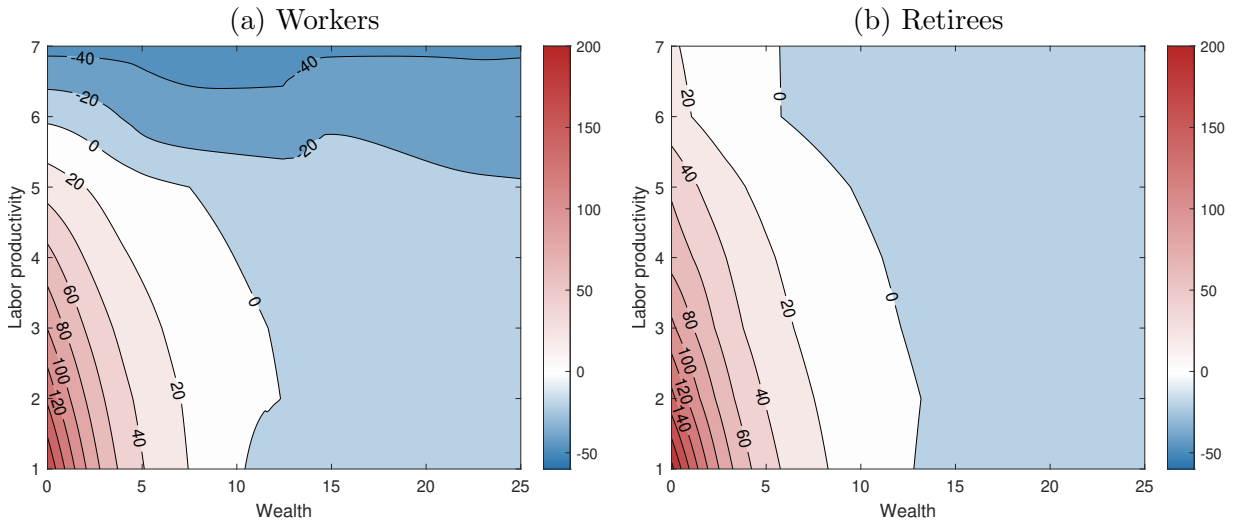
just over 2.2 percent in the initial steady state. To finance these large outlays, p_{SP} places high tax rates on consumption and capital income, 51.2 and 60.0 percent, respectively. At 60.0 percent, τ_y is also very high, but progressivity, ν_y, is higher than the benchmark but interior to the grid at 0.222 so low-income workers still pay little in labor taxes.

Enacting the optimal policy produces enormous average welfare gains (measured in consumption equivalent). When averaged over the initial living households, the welfare gain is 81.0 percent, and it is 87.8 percent when calculated "under the veil of ignorance."[13] Figure 8 plots the welfare gains induced by p_{SP} for workers and retirees along the productivity and wealth dimensions, aggregated at the skill level. The range of wealth depicted covers more than 99 percent of the mass of households.

One can observe that for both demographic categories, households with low wealth and low productivity achieve substantial welfare gains, with the vast majority in both graphs - in terms of mass in the distribution - having positive and large gains from the reform. A larger span of the retirees are better off at the optimal policy, showing only moderate losses whenever they are in the top tail of the wealth distribution. Working-age "stepping stars" and "superstars" lose from the reform, since there is no way for them to avoid the new high tax rates. Retired stars, however, typically enjoy a moderate welfare gain because their income is much lower than that of their younger counterparts. Only the most wealth-rich among them lose from the optimal policy. Finally, workers with productivity above the median show some heterogeneity in their welfare depending on their wealth, a disagreement that was highlighted in Figure 7. Overall, 85.9 percent of households support the reform.

[13]The "behind the veil of ignorance" measure excludes certain household types from the calculation because a person cannot be born into those states. These excluded states are stepping-stars, superstars, and retirees.

Figure 8: Welfare gains (in percent).



Notes: The figure shows the welfare gains of the optimal policy p_{SP} in terms of consumption equivalent variation for workers and retirees along the wealth and productivity grids. The x-axis shows the numerical wealth levels, which are capped at the value of 10 to highlight the part of the distribution in which there is the most mass of households. More than 90 percent of households are located at the range depicted. The y-axis scales productivity levels from the lowest to the highest, number with 7 being a "superstar" household. Panel (a) shows the heat map for workers and panel (b) for retirees. Both panels are aggregated at the skill level, showing the results for both skilled and unskilled households weighted by their relative masses.

5.1 Progressivity

It is important to keep in mind that each element of the policy is related to the others. If we had optimized over only one tax instrument while holding the others fixed, we would have come to very different conclusions about optimal tax policy. We demonstrate this for progressivity. Figure 9 plots social welfare as a function of progressivity in different feasible tax environments.[14] In each panel, two of the other three tax parameters are held fixed at their initial steady-state values, while the final one is varied (shown by different curves).

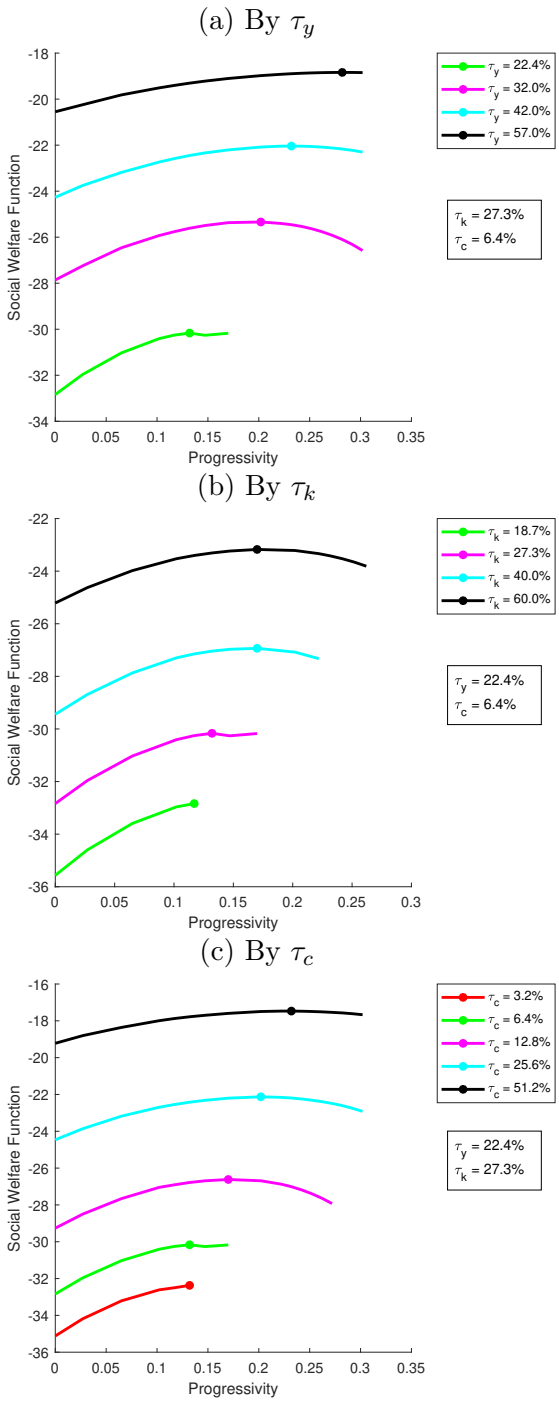
Panel (a) shows how optimal progressivity changes as the average level τ_y changes. The capital income tax rate and the consumption tax rate remain at their initial values of 27.3 percent and 6.4 percent, respectively. The "buying progressivity" concept is apparent: when labor taxes are low on average, the optimal schedule must be flattened in order to boost revenue and fund the transfer. As the overall tax level increases and the level of revenues rises, a greater degree of tax progressivity can be offered. Recall that lowest-income, lowest-

[14]Infeasible policies violate the non-negativity constraint on revenues. Their locations are suggested by curves that either terminate "early" or are not shown at all (for example, there is no line for $\tau_y = 5.0$ percent in panel (a)).

wealth agents have the highest weight in the social welfare function. Increasing the size of transfers by increasing τ_y unambiguously makes these agents better off and has the biggest impact on the social welfare function. However, unless progressivity is also adjusted, these gains in revenue come at the cost of adding to the tax burden of poor agents. With sufficiently high revenue, the social welfare function can be further optimized by shifting the tax burden even more from poor to rich households.

Optimal progressivity depends somewhat less on the level of capital income taxation (panel b). With τ_c and τ_y fixed at their initial values, the optimal ν_y shows almost no response to changes in τ_k. Meanwhile, consumption taxation sits between capital income taxes and average labor taxes. The trade-off between revenue and progressivity is visible in panel (c) but it is much more muted than in panel (a). The effect of consumption taxation on households' decisions shares aspects of both the labor and the capital income taxes. Like capital income taxation, consumption taxes distort intertemporal consumption/savings decisions though only in the first period. Like labor taxes, consumption taxes make leisure more attractive relative to consumption, reducing hours; however, this applies only to young households and the strength of this distortion wanes as household wealth rises.

Figure 9: Optimal progressivity.



Notes: The figure shows the values for the utilitarian social welfare function along the progressivity grid for each of the other three tax instruments. Each panel shows the function for different levels for a given tax rate with the other two fixed at their calibrated benchmark values. The top panel depicts it for the labor tax, the center panel for capital income tax, and the bottom panel for the consumption tax. Each line has a highlighted dot that indicates the welfare-maximizing level of progressivity at the given tax rate.

Optimal progressivity then depends critically on the rest of the tax environment, espe-

cially on the level of labor income taxation. This interdependence should be kept in mind when analyzing optimal tax progressivity in isolation from the overall level of taxation and redistribution since local optima could be far from the global optimum.

Similar analysis of the non-progressivity tax parameters offers only one additional insight: fixing a progressivity level and conditioning on any two of the other three tax parameters, welfare is maximized by the highest value of the remaining tax. Figures plotting these results can be viewed in Appendix D.1. This result follows from the strong motivation for transfers discussed above.

5.2 The Role of the Initial Wealth Distribution and Transitional Dynamics

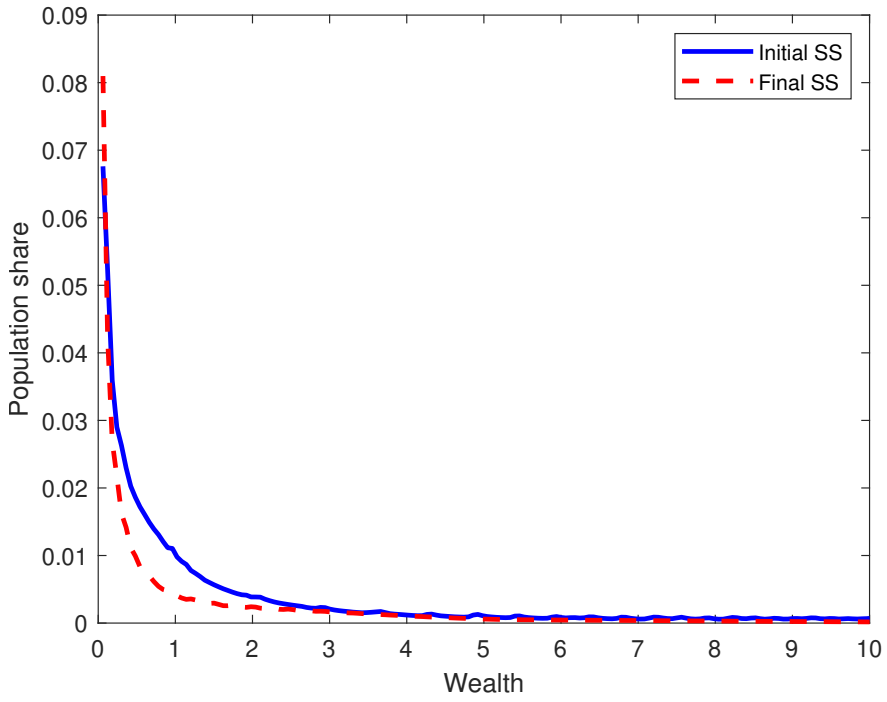
The optimal policy in this model features strikingly high tax rates on all sources, and consequently it leads to drastic reductions in economic activity over time. The central motivation behind the policy is that the revenues fund a large transfer that mitigates the stringent effects of market incompleteness on the poor. This can lead to a high-tax outcome for three reasons. First, the policy change is evaluated according to the preferences of households living in the initial steady state, an environment calibrated to match the high degree of income and wealth inequality in the US. The long right tail of accumulated past savings presents a tempting target for redistribution. Second, households face mortality risk, meaning that these decisive households likely live through only a small portion of the transition and so do not fully internalize the consequences of the reform on the capital stock. Third, policy is chosen "once-and-for-all." This disallows time-varying paths that may provide for redistribution of initial wealth inequality while also encouraging future capital accumulation as in Dyrda and Pedroni (2023). The benefit from initial redistribution then is entangled with that from providing a permanent higher level of social insurance against labor income risk.

We conduct two additional numerical experiments to better separate the effect of initial inequality and transitional costs on the optimal policy. First, we repeat our exercise but from the *final* steady state arising under the baseline optimal policy.[15] As a consequence of past high-tax policy, the wealth distribution inherited by these households is quite different. Relative to the initial distribution from the baseline, this "tax-and-transfer" wealth distribu-

[15]The number of feasible tax menus in $\mathcal{P}$ for this exercise is 1335.

tion has a 54.9 percent lower mean level of wealth and a far greater percentage of households that are borrowing constrained (55.5 vs 26.5). It is also severely compressed (Figure 10) so there is far less wealth to redistribute from the tail.

Figure 10: Initial and final densities of wealth.



Notes: The figure shows the initial and final steady-state densities along the wealth grid. The density was derived from the distribution by aggregating across all dimensions for all households. The density is truncated at the zero wealth level for visualization purposes and only shows households that are above the borrowing constraint. The x-axis shows the numerical wealth levels, which are capped at the value of 10 to highlight the part of the distribution in which there is the most mass of households.

One can think of this exercise as either solving the social planner's problem with fewer initial resources or, from the political economy perspective, as a surprise "re-vote" where households living through the consequences of their ancestors' redistribution are offered the chance to remake policy according to their wishes. As we will see below, we find that the social planner's policy is also the voting equilibrium policy, and, perhaps surprisingly, it is nearly identical to the "tax-and-transfer" optimal menu, which, apart from being quite more progressive, is qualitatively similar to the optimal policy found in the baseline experiment.

That high taxes are still chosen in this low-wealth environment undermines the argument that our baseline result was purely a consequence of initial wealth conditions combined with the assumption that the government can commit to a tax policy. Instead it reinforces the

intense preference for greater insurance. The new policy's higher degree of progressivity aligns well with this reasoning, as it follows mainly as a reflection of the higher tax elasticity of transfers. In the tax-transfer steady state, the economy starts at a point of increased equality in which there is substantially less capital available, and, hence a smaller tax base, requiring a more progressive labor income taxation to extract revenues. This allows the planner conducting the "revote" in the distant future to keep sustaining a high transfer-to-GDP ratio, though one that is now nearly 10 percentage points smaller than the one achieved by imposing the stringent optimal regime to the initial steady state.

While strongly suggestive, this experiment on its own is nevertheless not fully convincing. Policies with lower taxes that lead to steady states with higher capital also necessitate that households forgo consumption in the transition to build up that capital. It is possible then that households in the tax-and-transfer steady state would like to exit it for more prosperous outcomes but the transition costs imposed by the low-wealth initial condition dissuade them.

To investigate this further, we solve for the optimal policy that maximizes average steady-state welfare.[16] With transition costs removed from consideration, the optimal policy is $\{57.0\%, 30.2\%, 0.0\%, 51.2\%\}$ with a transfer-to-GDP ratio of 43.0 percent.

Table 5: Summary of optimal policies.

	τ_y	ν_y	τ_k	τ_c	Υ/Y
Baseline	57.0%	22.2%	60.0%	51.2%	61.1%
Tax-transfer	57.0%	27.2%	60.0%	51.2%	52.4%
Steady state only	57.0%	30.2%	0.0%	51.2%	43.0%

Notes: The table shows the optimal tax policy along with the present discounted value of transfers to GDP associated with each economy. "Baseline" evaluates policies from the calibrated initial steady state. "Tax-transfer" evaluates from the final steady state arising under the optimal policy in the baseline. "Steady state only" ignores transitional dynamics and maximizes average steady-state welfare.

The result is shown in Table 5 and its interpretation is straightforward. Households still want a lot of insurance. Consumption taxation is still at its highest allowable value.

[16] One could think of this as a "behind-the-veil-of-ignorance" exercise where a household chooses which steady state to live in, but does not know the household state it will start with. Optimal policy then maximizes expected utility in the steady state.

Labor income tax rates are at a corner and the progressivity of the schedule is now at its upper limit. High progressivity can co-exist with very large transfers because capital income taxation has been eliminated. The economy has a very high level of capital, which boosts wages and labor income, particularly among high earners.[17] The greater level of aggregate economic activity fostered by eliminating taxes on capital also expands the consumption tax base. This result highlights the importance of making a consumption tax available in the tax menu.

5.3 Intergenerational Welfare Effects

Up to this point, the analysis has focused on the welfare of households that are alive in the initial period of reform. Focusing only on the living is a natural perspective if one thinks of the central thought experiment as a study of the current demand for tax reform and redistribution by those facing an environment with a high degree of risk and heavily concentrated pools of resources to target. This focus imposes immense costs, however, not only on wealthy households in the initial period, but also on future generations of households yet to be born. The optimal policy depletes the capital stock, trading away future production for an immediate feast, which raises the question: "What if the social planner valued the well-being of future cohorts?" To answer this question, we conduct the "behind the veil of ignorance" experiment for every cohort along the transition path. As before, the household does not know what productivity and wealth it will be born with, but it does know the period of its birth and the wealth distribution from which it will draw that initial state.

Figure 11 plots the tax system that would maximize the average welfare of the newborns in each period *had that tax system been put in place from the first period onward*. This policy is a forward-looking behind the veil calculation: What distribution would each household want to be born into, understanding that their initial wealth will be chosen according to it?

All future cohorts agree (on average) on two things. First, the average tax on labor income should be very high, as should the consumption tax, reflecting a common demand for redistribution. This arises from the fact that the fundamental sources of risk in the economy, specifically persistent idiosyncratic shocks to productivity, exogenous borrowing

[17]The result of this dynamic shift of τ_k is also seen in Dyrda and Pedroni (2023), where the tax starts at very high levels and then declines as the government switches to labor income taxes. This substitution is a common result in Ramsey taxation as discussed in Conesa et al. (2009) and others.

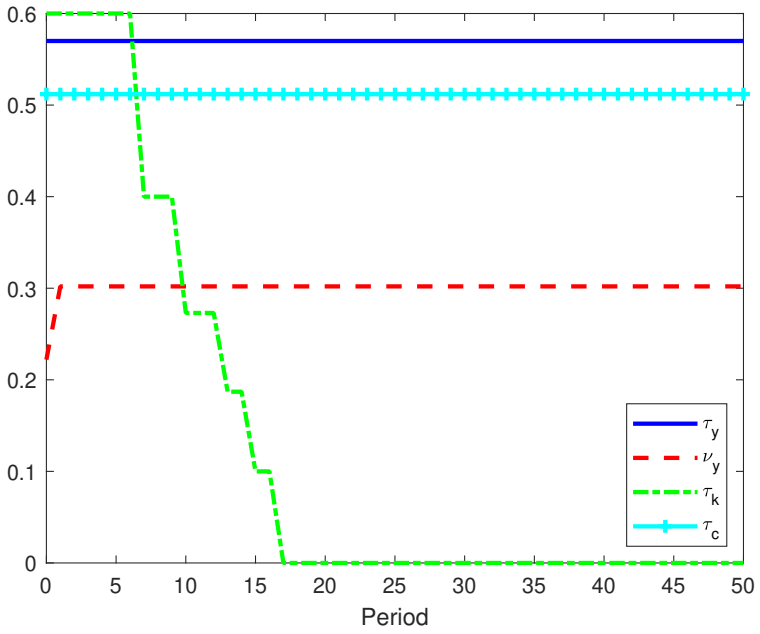
constraints, and mortality, are permanent.

Second, they agree that labor taxation should be highly progressive. This result differs from the preference of the initial living who want a less progressive schedule. As shown in Figure 2, high progressivity reduces transfers by strongly discouraging labor supply among the highly productive and the wealthy. For future cohorts, however, things are different; since all of the high-tax policies that these cohorts favor inevitably diminish the right tail over time, the wealth effect on high-income households will have weakened by the time these cohorts are born. In this way, for these later groups the transfer is less elastic with respect to progressivity and so it benefits them to tax high earners aggressively.

The most striking source of intergenerational conflict concerns how much to tax capital income. The initial living and the first six waves of newborns, who live in a world that still has a thick right wealth tail, benefit from targeting the wealthy. As a cohort's birth period is extended, however, the preference to tax the tail, which will have all but disappeared by then, gives way to a desire for low capital income taxes in order to encourage capital accumulation by earlier generations. All households born after period 15 of the transition favor eliminating capital income taxation entirely.[18] Figure 11 also illustrates that the optimal policy would be mostly unchanged if the utilitarian social welfare were augmented to include social discounting (Farhi and Werning, 2007). Only the capital income tax would be affected, and it would decrease as the relative planner's weight on future cohorts increased.

[18]This result echoes the well-known result from representative agent frameworks Chamley (1986) and Judd (1985) that capital income should not be taxed in the long run.

Figure 11: Most preferred initial tax reform by cohort.



Notes: The figure plots the preferred permanent tax reform of every cohort of young households born along the first 50 periods of transition. Specifically, for each cohort, the figure shows the tax combination that maximizes its average welfare *assuming that policy had been adopted from the initial period onward*. The x-axis displays the number of periods after the initial period in which the cohort arrives. Zero corresponds to the initial young households.

Given that future generations disagree strongly about the ideal level of capital taxation in the time-0 reform and given the severe consequences of that reform on the wealth distribution they inherit, one might wonder if some cohorts in the distant future would prefer that initial households had decided against reforming taxes. To answer this question, we compare expected newborn welfare for each cohort under both the optimal policy and under the initial policy from the calibrated benchmark economy (i.e., no reform). All cohorts strongly prefer the optimal policy over the initial one despite its high level of capital taxation. Put differently, given that these cohorts must live with the consequences of their ancestors' one-and-for-all reform, they prefer a high tax and transfer economy to the benchmark economy. Additionally, we also compare the optimal policy to the steady state welfare maximizing one, which has no capital income taxation. By definition, the later policy must be preferred by all cohorts born far enough in the future. We find that the optimal policy returns higher expected welfare for the first nine cohorts only. These results are displayed in Figure 16 in Appendix D.2.

5.4 Different Fiscal Regimes

In our benchmark economy we assumed that the government commits to spending the level G and to issuing and servicing the fixed value of debt B that are implicitly defined, respectively, by the calibration to 18 percent of output and clearing of the government budget constraint in the initial steady state. This type of fiscal regime is akin to the government committing to nominal values of spending and debt that were chosen before any fiscal reform. Since under the optimal tax policy, p_{SP}, the capital stock, along with labor and other aggregates, is heavily depleted along the transition to the new steady state, it becomes incrementally more expensive for the government to sustain fixed values of spending and debt. The cost is even higher in an economy where the lump-sum transfer rises above 60 percent of output.

In order to understand the effect of committing to a level or a share of spending and debt in the optimal choice of the tax menu, we consider three distinct fiscal regimes, defined by varying the commitment to either the level or the initial output share of G or B in the baseline economy: (i) "Fiscal Regime 1," in which spending remains fixed at the level in the benchmark economy's steady state and remains constant throughout the entire transition, i.e., $G_t = \bar{G}$, $\forall t$, but the government commits to the initial ratio of debt over GDP throughout the transition, i.e., $B_t = b_0 Y_t$, where b_0 is the share in the benchmark economy's steady state; (ii) "Fiscal Regime 2," where the government keeps the value of debt fixed as $B_t = \bar{B}$ and commits to the share of public spending $G_t = g_0 Y_t$, where, similarly, g_0 is the share of government expenditure calibrated in the benchmark economy; and (iii) "Fiscal Regime 3," in which both debt and expenditure are held at the constant shares throughout the transition, with $B_t = b_0 Y_t$ and $G_t = g_0 Y_t$, $\forall t$. We then re-run our main numerical experiment and reconstruct the menu of fiscal policies $\mathcal{P}$ for each of these three regimes and compare them to the benchmark economy.[19] Table 6 shows the optimal tax menu for each of these regimes.

[19] The number of feasible cases of $\mathcal{P}$ for each of these regimes is, respectively, 2556, 2584, and 2606.

Table 6: Summary of optimal policies under different fiscal regimes.

	τ_y	ν_y	τ_k	τ_c	B/Y	G/Y	Υ/Y
Baseline	57.0%	22.2%	60.0%	51.2%	81.1%	22.8%	61.1%
Fiscal Regime 1	57.0%	23.2%	60.0%	51.2%	64.1%	22.6%	61.6%
Fiscal Regime 2	57.0%	24.2%	60.0%	51.2%	83.4%	18.0%	68.1%
Fiscal Regime 3	57.0%	26.2%	60.0%	51.2%	64.1%	18.0%	68.2%

Notes: The table shows the optimal tax policy along with the present discounted value of transfers to GDP, government spending to GDP, and debt to GDP, associated with each fiscal regime. "Baseline" evaluates policies from the calibrated initial steady state. "Fiscal Regime 1" refers to the optimal policy with a fixed level of spending and a fixed share of debt-to-GDP. "Fiscal Regime 2" refers to the optimal policy with a fixed share of spending-to-GDP and a fixed level of debt. "Fiscal Regime 3" refers to the optimal policy with a fixed share of spending and debt to GDP.

Similarly as in previous exercises, the optimal policy assigns to the consumption tax, to the capital income tax, and to the parameter that governs the average labor tax the maximum allowed rates defined in each of their grids. In the same spirit of the previous results and exercises, the parameter governing the curvature of the progressive tax function, ν_y, is the only one for which the choice remains interior at the grid.

In Table 6 we observe again the concept of "buying progressivity" as we move down the rows from the baseline economy to the economy under "Fiscal Regime 3". In all of the regimes, and similarly to what was shown in Figure 1, capital is eaten along the transition, achieving a substantially lower level in each optimal steady state. Hence, the regime of the benchmark economy is relatively more expensive for the government to finance in the long run. The intuition here is simple: as we move toward cheaper regimes, i.e., regimes in which we commit to fixed fractions of output, the government obtains more fiscal room in the long term to raise progressivity and, consequently, the size of the lump-sum transfer relative to output, Υ/Y. From the benchmark economy to "Fiscal Regime 1" we fix B/Y, which is relatively more expensive than G/Y, hence making "Fiscal Regime 2" a bit cheaper than the latter and leaving "Fiscal Regime 3" the one in which both rates are fixed, as the cheapest one that yields the highest choice of ν_y.

6 Discussion - Head-to-Head Voting

The socially optimal levels of taxation and redistribution in the model economy are far greater than anything observed in the data. One possible explanation for this discrepancy is that social planners are not relevant. In practice, government policy is chosen through some political process, which naturally involves strategic interactions between self-interested parties and therefore can easily deliver suboptimal outcomes. With this in mind, we examine a second method for aggregating preferences in our model: majority voting.

We search for all tax policies that could arise from a sequence of head-to-head elections among all p in the tax space, $\mathcal{P}$. It is well-known that when the policy space is multidimensional (as it is here), the final outcome of these elections can depend on how elections are ordered in the sequence.[20] We follow the method described in Carroll et al. (2021) to identify political equilibria in our model economy. This method uses the initial distribution of households over X and the indirect utility functions, $\{V(\boldsymbol{p}; x)\} \forall p \in \mathcal{P}$, and conducts several successive refinements of $\mathcal{P}$. At the end of this process, the remaining set contains all the policies that could be an equilibrium even under strategic voting and agenda setting.[21] The benefit of this method is that it does not require any onerous additional structure to be imposed on either household preferences or the voting process. Should multiple political equilibria exist, this method will find them.[22]

Although we adopt this robust method for identifying political equilibria, we never find a set that is other than a singleton. That is, there is a unique policy, p_{Cond}, in $\mathcal{P}$ that defeats all other policies in any arrangement of head-to-head elections.[23] Carroll et al. (2021) provide some reasoning for why a unique equilibrium is not surprising in our environment. Essentially, it comes down to two features of the model. First, as shown in Section 4.2, there is a lot of disagreement among households over fiscal policy, and second, no single household type has a large weight. Taken together, these factors ensure that the model will not have just a few voting blocs that can be combined in multiple ways to reach a majority.

The political equilibrium under majority voting is nearly identical to the one that max-

[20] See Persson and Tabellini (2002) for a deeper discussion of multidimensional voting.

[21] For example, if one agent can sequence the pairwise competition she may be able to alter the outcome. Our method identifies any outcome that could win in some sequence of pairwise votes.

[22] As discussed in Carroll et al. (2021), the method relies on $\mathcal{P}$ being finite, which is satisfied by construction in our experiment.

[23] Such a policy is known as a "Condorcet" winner.

imizes social welfare in (13), with the exception of a stronger desire for progressivity, with $p_{Cond} = \{57.0\%, 26.2\%, 60.0\%, 51.2\%\}$. In this sense, our analysis revisits a classical argument highlighted by Aiyagari and Peled (1995) in which majority voting and utilitarian planner outcomes do not yield substantively different results. In that paper the similarity of the two equilibrium concepts hinges critically on agents' need for insurance. When the need for insurance is very low, the utilitarian planner will tax less than the majority voting outcome. As the desire for insurance increases, the utilitarian planner is *more responsive* than majority voting to this demand. If households' risk aversion and idiosyncratic shock volatility fall in ranges consistent with the data, the policies from both mechanisms are roughly equal.

The same principle applies here. Under "one man, one vote," the political weight of any household type is just equal to its population share in the model economy. Relative to this benchmark, a utilitarian social planner takes into account not only the population share but also the marginal utility of each type. The initially wealth-poor, particularly those with low productivity draws, have very high present and expected future marginal consumption. The boost to social welfare from shifting resources toward them far outweighs the loss suffered by the initially rich. Thus, the social planner further shifts policy toward the poor by giving an even larger weight to the poor.

Nonetheless, the relatively large population share of young households helps in explaining the stronger desire for progressivity in the voting outcome, as the groups that would oppose it are mainly the "stars" and the retirees. Due to their smaller population share and preference for lower progressivity, as shown in Table 4, these groups are outvoted in the head-to-head protocol, being unable to leverage as much influence, albeit small, as they had in the social planner's scheme. The steep welfare loss suffered by the stars (as shown in Figure 8), combined with the loss of revenues for the lump-sum transfer much desired by the retirees, when compared to their voting power, yields a stronger weight in the utilitarian SWF that is applied by the planner to this group.

7 Conclusion

We have explored the optimal tax-and-transfer policy in a model environment with rich household heterogeneity and where the government has many tools for raising tax revenue.

The optimal policy in the model places very high tax rates on capital income and on consumption. Labor income taxes feature high top tax rates but also a progressive schedule that is higher than what is currently observed in the US code. Greater tax progressivity is welfare reducing because it trades away a higher transfer in order to extend tax reductions to households with moderate marginal utilities of consumption. When a policy is instead decided by majority voting in the model economy, the outcome is identical to the social planner's choice for all tax instruments except progressivity, which is slightly higher.

The fiscal policy predicted by the model is quite different than any seen in the data. Making the two more consistent would require shifting policy more in the direction of that preferred by more educated and wealthier households. In the follow-up short manuscript Carroll et al. (2024), we provide empirical evidence that political identity and ideology are strong predictors of people's preferences over tax, transfers, and redistribution. We estimate that voters seem to choose candidates with tax platforms that do not maximize their economic interests but better satisfy them along other dimensions unrelated to fiscal policy. It is possible, in light of our model results and empirical evidence, that in a political environment in which parties choose candidates and seek to maximize vote shares, a candidate may not emerge who satisfies both economic and non-economic interests for a substantial number of voters.

Furthermore, we have assumed that production is Cobb-Douglas in our model economy so that the elasticity of substitution between factors is always equal to one. Although Cobb-Douglas production is a common specification in the literature and perhaps a sufficient approximation to the aggregate production function in developed countries, it is doubtful that this assumption is reasonable in an economy with such a low level of capital as the one resulting from the equilibrium policy. It seems likely that the elasticity of substitution would decline substantially given such a large imbalance of input factors, and that this decline would reduce the incentives to impose very high taxes on the model economy's resources to fund large transfers. We plan to explore this line of inquiry in the future.

References

Abraham, Arpad, Pavel Brendler, and Eva Carceles (2023). "Optimal Income Redistribution." Technical report. URL https://drive.google.com/file/d/1-gu_dg1L07MN05UsYyf_

M5paY4dZRa43/view.

Acikgoz, Omer, Marcus Hagedorn, Hans Aasnes Holter, and Yikai Wang (2018). "The Optimum Quantity of Capital and Debt." Technical report. doi:10.2139/ssrn.3185088.

Aguiar, Mark and Mark Bils (2015). "Has consumption inequality mirrored income inequality?" *American Economic Review*, 105(9), pp. 2725–2756. doi:10.1257/aer.20120599.

Aiyagari, Rao (1994). "Uninsured Idiosyncratic Risk and Aggregate Saving." *The Quarterly Journal of Economics*, 109(3), pp. 659–684. doi:10.2307/2118417.

Aiyagari, Rao and Dan Peled (1995). "Social insurance and taxation under sequential majority voting and utilitarian regimes." *Journal of Economic Dynamics and Control*, 19(8), pp. 1511–1528. doi:10.1016/0165-1889(94)00840-E.

Bachmann, Rüdiger and Jinhui H. Bai (2013). "Politico-economic inequality and the comovement of government purchases." *Review of Economic Dynamics*, 16(4), pp. 565–580. doi:10.1016/j.red.2012.09.008.

Bakış, Ozan, Barış Kaymak, and Markus Poschke (2015). "Transitional dynamics and the optimal progressivity of income redistribution." *Review of Economic Dynamics*, 18(3), pp. 679–693. doi:10.1016/j.red.2014.08.004.

Benabou, Roland (2002). "Tax and Education Policy in a Heterogeneous-Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency?" *Econometrica*, 70(2), pp. 481–517. doi:10.1111/1468-0262.00293.

Bewley, Truman (1986). "Stationary monetary equilibrium with a continuum of independently fluctuating consumers." In Werner Hildenbrand and Andreu Mas-Collel, editors, *Contributions to Mathematical Economics in Honor of Gérard Debreu*, volume 79. North-Holland Amsterdam.

Blanchard, Olivier J. (1985). "Debt, Deficits, and Finite Horizons." *Journal of Political Economy*, 93(2), pp. 223–247. doi:10.1086/261297.

Boar, Corina and Virgiliu Midrigan (2022). "Efficient redistribution." *Journal of Monetary Economics*, 131, pp. 78–91. doi:10.1016/j.jmoneco.2022.07.003.

Carey, David and Josette Rabesona (2003). "Tax Ratios on Labour and Capital Income and on Consumption." *OECD Economic Studies*, 2002(2), pp. 129–174. doi:10.1787/eco_studies-v2002-art11-en.

Carroll, Daniel and Sewon Hur (2023). "On the Distributional Effects of International Tariffs." *International Economic Review*, 64(4), pp. 1311–1346. doi:10.1111/iere.12648. URL https://onlinelibrary.wiley.com/doi/10.1111/iere.12648.

Carroll, Daniel R., Jim Dolmas, and Eric R. Young (2021). "The politics of flat taxes." *Review of Economic Dynamics*, 39, pp. 174–201. doi:10.1016/j.red.2020.06.016.

Carroll, Daniel R., André Victor D. Luduvice, and Eric R. Young (2024). "A Note on Aggregating Preferences for Redistribution." Working paper (Federal Reserve Bank of Cleveland). doi:10.26509/frbc-wp-202427. URL https://www.clevelandfed.org/publications/working-paper/wp-2427-aggregating-preferences-for-redistribution, series: Working paper (Federal Reserve Bank of Cleveland).

Castañeda, Ana, Javier Díaz-Giménez, and José-Víctor Ríos-Rull (2003). "Accounting for the U.S. Earnings and Wealth Inequality." *Journal of Political Economy*, 111(4), pp. 818–857. doi:10.1086/375382. URL https://www.journals.uchicago.edu/doi/10.1086/375382.

CBO (2019). "Mandatory spending in 2018: An infographic." Technical report, Congressional Budget Office, Washington, D.C.

Chamley, Christophe (1986). "Optimal taxation of capital income in general equilibrium with infinite lives." *Econometrica*, pp. 607–622.

Chang, Bo Hyun, Yongsung Chang, and Sun-Bin Kim (2018). "Pareto weights in practice: A quantitative analysis across 32 OECD countries." *Review of Economic Dynamics*, 28, pp. 181–204. doi:10.1016/j.red.2017.08.002.

Conesa, Juan Carlos and Carlos Garriga (2003). "The status quo problem in social security reform." *Macroeconomic Dynamics*, 7, pp. 691–710.

Conesa, Juan Carlos, Sagiri Kitao, and Dirk Krueger (2009). "Taxing Capital? Not a Bad Idea after All!" *American Economic Review*, 99(1), pp. 25–48. doi:10.1257/aer.99.1.25.

Conesa, Juan Carlos, Bo Li, and Qian Li (2023). "A quantitative evaluation of universal basic income." *Journal of Public Economics*, 223, p. 104,881. doi:10.1016/j.jpubeco.2023.104881.

Corbae, Dean, Pablo D'Erasmo, and Burhanettin Kuruscu (2009). "Politico-economic consequences of rising wage inequality." *Journal of Monetary Economics*, 56(1), pp. 43–61. doi:10.1016/j.jmoneco.2008.08.012.

Daruich, Diego and Raquel Fernandez (2024). "Universal basic income: A dynamic assessment." *American Economic Review*, 114(1), pp. 38–88. doi:10.1257/aer.20221099.

de Souza, Gustavo (2022). "On the Political and Economic Determinants of Redistribution: Economic Gains, Ideological Gains, or Institutions?" Working paper 2022-47, Federal Reserve Bank of Chicago. doi:10.21033/wp-2022-47.

Dyrda, Sebastian and Marcelo Pedroni (2023). "Optimal Fiscal Policy in a Model with Uninsurable Idiosyncratic Income Risk." *The Review of Economic Studies*, 90(2), pp. 744–780. doi:10.1093/restud/rdac031.

Farhi, Emmanuel and Ivan Werning (2007). "Inequality and social discounting." *Journal of political economy*, 115(3), pp. 365–402.

Ferraro, Domenico and Vytautas Valaitis (2024). "Consumption Quality and Employment Across the Wealth Distribution." *The Review of Economic Studies*, p. rdae052. doi:10.1093/restud/rdae052.

Ferriere, Axelle, Philipp Grübener, Gaston Navarro, and Oliko Vardishvili (2022). "On the optimal design of transfers and income-tax progressivity." International Finance Discussion Papers 1350, Board of Governors of the Federal Reserve System. doi:10.17016/IFDP.2022.1350.

Guner, Nezih, Remzi Kaygusuz, and Gustavo Ventura (2023a). "Rethinking the welfare state." *Econometrica*, 91(6), pp. 2261–2294. doi:10.3982/ECTA19921.

Guner, Nezih, Martin Lopez-Daneri, and Gustavo Ventura (2016). "Heterogeneity and Government revenues: Higher taxes at the top?" *Journal of Monetary Economics*, 80, pp. 69–85. doi:10.1016/j.jmoneco.2016.05.002.

Guner, Nezih, Martin Lopez-Daneri, and Gustavo Ventura (2023b). "The looming fiscal reckoning: Tax distortions, top earners, and revenues." *Review of Economic Dynamics*, 50, pp. 146–170. doi:10.1016/j.red.2023.07.003.

Heathcote, Jonathan, Fabrizio Perri, and Giovanni L. Violante (2010). "Unequal we stand: An empirical analysis of economic inequality in the United States, 1967–2006." *Review of Economic Dynamics*, 13(1), pp. 15–51. doi:10.1016/j.red.2009.10.010.

Heathcote, Jonathan, Fabrizio Perri, Giovanni L. Violante, and Lichen Zhang (2023). "More unequal we stand? Inequality dynamics in the United States, 1967–2021." *Review of Economic Dynamics*, 50, pp. 235–266. doi:10.1016/j.red.2023.07.014.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante (2017). "Optimal Tax Progressivity: An Analytical Framework." *The Quarterly Journal of Economics*, 132(4), pp. 1693–1754. doi:10.1093/qje/qjx018.

Heathcote, Jonathan and Hitoshi Tsujiyama (2021). "Optimal Income Taxation: Mirrlees Meets Ramsey." *Journal of Political Economy*, 129(11), pp. 3141–3184. doi:10.1086/715851.

Holter, Hans A., Dirk Krueger, and Serhiy Stepanchuk (2019). "How do tax progressivity and household heterogeneity affect Laffer curves?" *Quantitative Economics*, 10(4), pp. 1317–1356. doi:10.3982/QE653.

Huggett, Mark (1993). "The risk-free rate in heterogeneous-agent incomplete-insurance economies." *Journal of Economic Dynamics and Control*, 17(5-6), pp. 953–969. doi:10.1016/0165-1889(93)90024-M.

Huggett, Mark and Juan Carlos Parra (2010). "How Well Does the US Social Insurance System Provide Social Insurance?" *Journal of Political Economy*, 118(1), pp. 76–112. doi:10.1086/651513. URL https://www.journals.uchicago.edu/doi/10.1086/651513, tex.ids= huggettHowWellDoes2010a.

Imrohoroğlu, Ayşe (1989). "Cost of Business Cycles with Indivisibilities and Liquidity Constraints." *Journal of Political Economy*, 97(6), pp. 1364–1383. doi:10.1086/261658.

Imrohoroğlu, Ayşe, Cagri Kumru, and Arm Nakornthab (2018). "Revisiting Tax on Top Income." Working Papers in Economics and Econometrics 660, Australian National University. URL https://cbe.anu.edu.au/researchpapers/econ/wp660.pdf.

Jones, John Bailey and Yue Li (2023). "Social Security reform with heterogeneous mortality." *Review of Economic Dynamics*, 48, pp. 320–344. doi:10.1016/j.red.2022.06.003.

Judd, Kenneth L (1985). "Redistributive taxation in a simple perfect foresight model." *Journal of Public Economics*, 28(1), pp. 59–83.

Kaymak, Barış, David Leung, and Markus Poschke (2023). "Accounting for Wealth Concentration in the United States." *Conditionally Accepted, American Economic Journal: Macroeconomics*. doi:10.26509/frbc-wp-202228.

Kaymak, Barış and Markus Poschke (2016). "The evolution of wealth inequality over half a century: The role of taxes, transfers and technology." *Journal of Monetary Economics*, 77, pp. 1–25. doi:10.1016/j.jmoneco.2015.10.004. URL https://linkinghub.elsevier.com/retrieve/pii/S0304393215001233.

Kina, Özlem, Ctirad Slavík, and Hakki Yazici (2024). "Redistributive capital taxation revisited." *American Economic Journal: Macroeconomics*, 16(2), pp. 182–216. doi:10.1257/mac.20200395.

Kindermann, Fabian and Dirk Krueger (2022). "High Marginal Tax Rates on the Top 1 Percent? Lessons from a Life-Cycle Model with Idiosyncratic Income Risk." *American Economic Journal: Macroeconomics*, 14(2), pp. 319–366. doi:10.1257/mac.20150170.

Krueger, Dirk, Kurt E. Mitman, and Fabrizio Perri (2016). "Macroeconomics and Household Heterogeneity." In *Handbook of Macroeconomics*, volume 2, pp. 843–921. Elsevier. doi:10.1016/bs.hesmac.2016.04.003.

Krueger, Dirk and Fabrizio Perri (2006). "Does income inequality lead to consumption inequality? Evidence and theory." *The Review of Economic Studies*, 73(1), pp. 163–193. doi:10.1111/j.1467-937X.2006.00373.x.

Krusell, Per, Vincenzo Quadrini, and José-Víctor Ríos-Rull (1997). "Politico-economic equilibrium and economic growth." *Journal of Economic Dynamics and Control*, 21(1), pp. 243–272. doi:10.1016/0165-1889(95)00931-0.

Kuhn, Moritz and José-Víctor Ríos-Rull (2015). "2013 Update on the US Earnings, Income, and Wealth Distributional Facts: A View from Macroeconomics." Technical report. URL https://www.sas.upenn.edu/~vr0j/papers/QR5_revised_final_thanks.pdf.

Lopez-Daneri, Martin (2016). "NIT picking: The macroeconomic effects of a Negative Income Tax." *Journal of Economic Dynamics and Control*, 68, pp. 1–16. doi:10.1016/j.jedc.2016.04.008.

Luduvice, André Victor Doherty (2024). "The macroeconomic effects of universal basic income programs." *Journal of Monetary Economics*, p. 103615. doi:10.1016/j.jmoneco.2024.103615.

Macnamara, Patrick, Myroslav Pidkuyko, and Raffaele Rossi (2023). "Taxing consumption in unequal economies." *SSRN Electronic Journal*. doi:10.2139/ssrn.4632460. URL https://www.ssrn.com/abstract=4632460.

McDaniel, Cara (2007). "Average Tax Rates on Consumption, Investment, Labor, and Capital in the OECD 1950-2003." Technical report. URL https://citeseerx.ist.psu.edu/document?repid=rep1&type=pdf&doi=c7f5ff20102c4434e653ba2af3f2acb114daf1da.

OMB (2023). "White House's Office of Management and Budget, Historical Tables." https://www.whitehouse.gov/omb/budget/historical-tables/.

Persson, Torsten and Guido Tabellini (2002). *Political Economics: Explaining Economic Policy*. MIT Press.

Tax Policy Center (2023). "Historical highest marginal income tax rates." https://www.taxpolicycenter.org/statistics/historical-highest-marginal-income-tax-rates.

Trabandt, Mathias and Harald Uhlig (2011). "The Laffer curve revisited." *Journal of Monetary Economics*, 58(4), pp. 305–327. doi:10.1016/j.jmoneco.2011.07.003.

Wu, Chunzan (2021). "More unequal income but less progressive taxation." *Journal of Monetary Economics*, 117, pp. 949–968. doi:10.1016/j.jmoneco.2020.07.005.

Yaari, Menahem E. (1965). "Uncertain Lifetime, Life Insurance, and the Theory of the Consumer." *The Review of Economic Studies*, 32(2), pp. 137–150. doi:10.2307/2296058.

Appendix

A Discounting with Stochastic Aging

Since our economy has stochastic aging, in order to properly measure the value of aggregates and welfare along the transition, we have to adjust their value and discount it to the required periods. We start by defining the sequential value of a given flow value, f_t, for a given worker at time t. In order to spare notation and allow the algebra to be representative for models with similar structure, we will omit, without loss of generality, the elements of the state space x and will only use time arguments. Hence, the value function for a retired worker at time t can be represented as:

$$V^R(t) = f_t + \beta(1-\psi_d)f_{t+1} + \cdots = \sum_{i=t}^{\infty} \left[\beta(1-\psi_d)\right]^{i-t} f_i \tag{14}$$

For a young worker, we can use the value in equation (14) to define her value function at time t:

$$V^W(t) = f_t + \beta\left[(1-\psi_a)V^W(t+1) + \psi_a V^R(t+1)\right] \tag{15}$$

Evaluating the value in (15) at $t+1$ and substituting it into time t, we have that:

$$V^W = f_t + \beta\left\{(1-\psi_a)\left[f_{t+1} + \beta\left[(1-\psi_a)V^W(t+2) + \psi_a V^r(t+2)\right]\right] + \psi_a V^R(t+1)\right\} \tag{16}$$

Rearranging the equation above yields:

$$V^W(t) = f_t + \beta(1-\psi_a)f_{t+1} + \beta^2(1-\psi_a)^2 V^W(t+2) + \beta(1-\psi_a)^2\psi_a V^R(t+2) + \beta\psi_a V^R(t+1) \tag{17}$$

Hence, iterating the previous step until a certain finite period of time T, we can compute the present discounted value for the worker at time t:

$$V^W(t) = f_t + \sum_{j=1}^{T} \left[\beta(1-\psi_a)\right]^j f_{t+j} + \sum_{j=1}^{T} \beta^j (1-\psi_a)^{j-1} \psi_a V^R(t+j) + \left[\beta(1-\psi_a)\right]^T V^W(t+T) \tag{18}$$

We can then also compute the steady-state values of the flows for both household types. From the retired household value in (14), if we start at $t = 0$, we have that:

$$V^{R,SS} = \frac{f^{SS}}{1-\beta(1-\psi_d)} \tag{19}$$

For the worker, we can take the limit and evaluate (18) at $T \to \infty$, starting at $j = 0$ to obtain:

$$V^W(t) = \sum_{j=0}^{\infty} \left[\beta(1-\psi_a)\right]^j f_{t+j} + \beta\psi_a \sum_{j=0}^{\infty} \left[\beta(1-\psi_a)\right]^j V^R(t+j+1) \tag{20}$$

where we used the fact that $\lim_{T\to\infty} \beta^T (1-\psi_a)^T V^W(t+T) = 0$, since V^W is bounded above.

Hence, in the steady state, we have that:

$$V^W = \frac{f^{SS} + \beta\psi_a V^{R,SS}}{1-\beta(1-\psi_a)} \tag{21}$$

We can then construct a simple algorithm to compute the present discounted value of a given flow f_t at the transition:

1. Start at T, with T large representing the machine equivalent of infinity, and approximate the time limit for the transition. From the steady-state value, $V^{R,SS}$, we know that $\frac{f^{SS}}{1-\beta(1-\psi_d)}$, where f^{SS} is already computed for the steady-state economy.

2. Define two auxiliary functions $A_1(t)$ and $A_2(t)$, which can be defined at time T as

$$A_1(T) \equiv f_T + \beta\left[(1-\psi_a)V^{W,SS} + \psi_a V^{R,SS}\right] \tag{22}$$

$$A_2(T) \equiv f_T + \beta(1-\psi_d)V^{R,SS} \tag{23}$$

3. Hence, in $T-1$, we can compute:

$$A_1(T-1) = f_{T-1} + \beta\left[(1-\psi_a)A_1(T) + \psi_a A_2(T)\right] \tag{24}$$

$$A_2(T-1) = f_{T-1} + \beta(1-\psi_d)A_2(T) \tag{25}$$

4. Iterate backward to $t = 1$ to find the values at the enacted period of the transition:

$$A_1(1) = f_1 + \beta\left[(1-\psi_a)A_1(2) + \psi_a A_2(2)\right] \tag{26}$$

$$A_2(1) = f_1 + \beta(1-\psi_d)A_2(2) \tag{27}$$

B Computation of the Model

We solve the model in several steps. First, for each $p = [\tau_y, \nu_y, \tau_k, \tau_c]$ from the grids in Section 4, we solve for the recursive competitive equilibrium (RCE) detailed in Section 2.2. This involves first solving for the steady state under p and then for the transition path back to the initial steady state. This also returns households' indirect utility, $V_x(p)$, from implementing reform p. Using these indirect utilities and the initial wealth distribution, Γ_1, we compute social welfare according to equation (13).

Solving for an RCE is done in the usual way. To find a steady state use the following steps:

1. Guess a rental rate r, a lump-sum transfer Υ and aggregate average earnings AE. Since the wage w can be expressed as a function of r from the firms' first-order conditions, households have all the information they need to solve their problem.

2. Solve the household problem given the guess at r, Υ, and AE.

3. Beginning with some initial wealth distribution, iterate on the distribution using the household decision rules. Repeat until the sup norm over the difference between any two consecutive distributions is less than a very small tolerance.

4. Use the converged wealth distribution and decision rules to check that, at r, aggregate capital supplied by the households equals the firm's demand, that Υ clears the government budget constraint and aggregate average earnings are consistent. If not, then update the guesses for r, Υ, and AE and repeat the steps above.

A transition path is solved in a similar way.

1. Assume that the transition is completed in T periods, guess *a sequence* of rental rates r_t, lump-sum transfers Υ_t, and average earnings AE_t.

2. Use $V_T(x)$ to solve the household problem in $T-1$. Along with $V_{T-1}(x)$, this also yields household decision rules in $T-1$, $g_{T-1}(x)$. Iterate backward to $t=1$, collecting the household decisions for all periods.

3. Starting at the initial steady-state distribution, Γ_1, use $g_1(x)$ to find Γ_2 and compute all time 1 aggregate variables. Repeat until the entire sequence of distributions from $1, ..., T$ has been found along with the associated sequences of capital supplies and government surpluses.

4. Check that the capital market and government budget constraint clear in every period. If not, update the guessed sequences using the values implied by the firm's first-order condition for capital demand and the government budget constraint in each period. As is customary, to better ensure convergence we use a dampening factor to update the guess slowly.

Once an RCE has been found for each p in $\mathcal{P}$, we compute the outcome under majority voting at the enacted period of the transition. We do this using the method detailed in Carroll et al. (2021). We use the following steps:

1. First, discard any policy that would be unanimously defeated by another policy in $\mathcal{P}$. The remaining policies form the Pareto set.

2. Next, reduce the Pareto set to the uncovered set by constructing the **adjacency matrix** M.

   - M is a square matrix of 0's and 1's with a dimension N, where N is the cardinality of $\mathcal{P}$.
   - Ordering the policies in $\mathcal{P}$ by $1, 2, ..., N$, $M(i, j)$ equals 1 if policy i defeats policy j in a head-to-head competition and 0 if not.

3. From M, the uncovered set can be found by computing the matrix

$$M^* = M^2 + M + I.$$

- Any policy i for which there exists a policy j with $m^*_{i,j} = 0$ is said to be covered by j.
- The uncovered set then consists of all the policies for which the corresponding rows of M^* contain only 1's.

4. If only a single policy satisfies this criterion, then it is the Condorcet winner policy. As stated in Section 6, we always find a Condorcet winner.

C Robustness: Expanding the Grids

Three of the rates in p_{SP} are at their upper bounds on the grid, suggesting that the planner would like to further raise those rates and redistribute even more. We have explored extending these upper limits by adding $\{70.0\%, 90.0\%\}$ to the τ_y, τ_k, and τ_c grids and then interacting these with the full grid for ν_y. This generates 144 new tax policies. Not all of these new policies, however, are feasible. This failure can arise because the combination of τ_y and the Social Security tax implies a marginal tax rate on earnings that exceeds 1. In other cases, the policy violates the non-negativity constraint on transfers. Finally, all of these high-tax policies produce wealth distributions that are close to degenerate as a very large of fraction of households hold zero assets. In this environment, the solution algorithm for finding market-clearing prices becomes highly unstable as small differences in the interest rate move a large mass of agents on or off the borrowing constraint.

Among the solutions that we could find, all produce higher average welfare and transfer-to-GDP ratio than the one we find in the main text, p_{SP}, with the highest welfare arising from $p = \{70.0\%.10.2\%, 70\%, 90.0\%, 66\%\}$. The biggest difference between this policy and p_{SP} is the degree of progressivity. In our main analysis, higher tax rates were associated with higher progressivity, while, in this robustness exercise, optimal progressivity is smaller than the initial calibrated value. While we view this extra-high-tax result as interesting, it is important to stress that p_{SP} is itself already yielding higher taxes than any fiscal policy we see implemented in practice, making this case even more counterfactual to the data. In this way, the qualitative nature of the solutions is similar even if the quantitative results differ from the pattern in our main exercise.

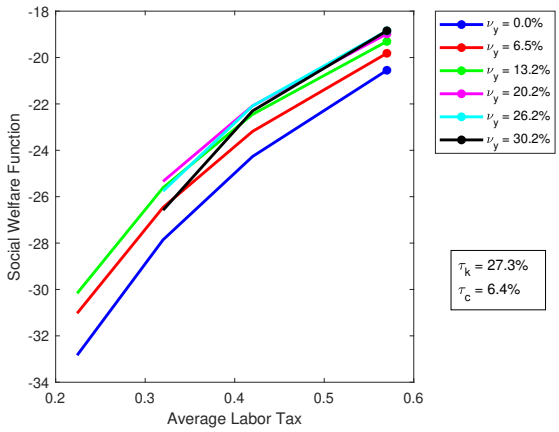
D Additional Figures

D.1 Optimal Tax Rates

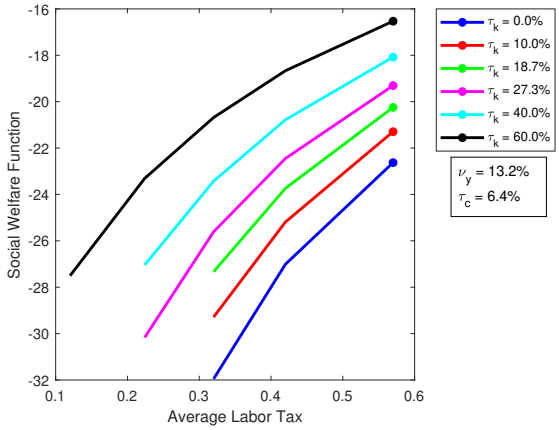
Figures 12-14 plot the relationships between the non-progressivity tax parameters in the model. In every case, social welfare is maximized when these parameters are at their highest feasible values.

Figure 12: Optimal τ_y

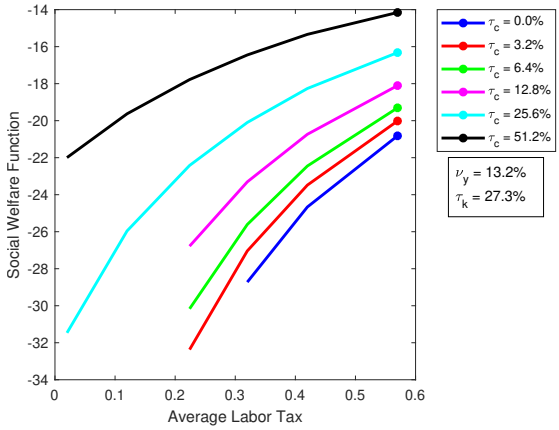
(a) By ν_y



(b) By τ_k



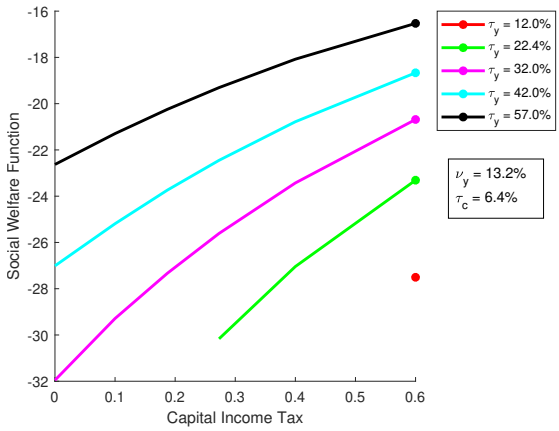
(c) By τ_c



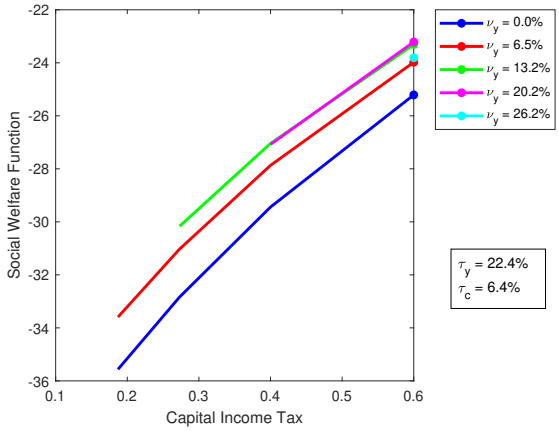
Notes: The figure shows the values for the utilitarian social welfare function along the average labor tax grid for each of the other three tax instruments. Each panel shows the function for different levels for a given tax rate with the other two fixed at their calibrated benchmark values. The top panel depicts it for the progressivity parameter, the center panel for capital income tax, and the bottom panel for the consumption tax. Each line has a highlighted dot that indicates the welfare-maximizing level of average labor tax at the given tax rate.

Figure 13: Optimal τ_k

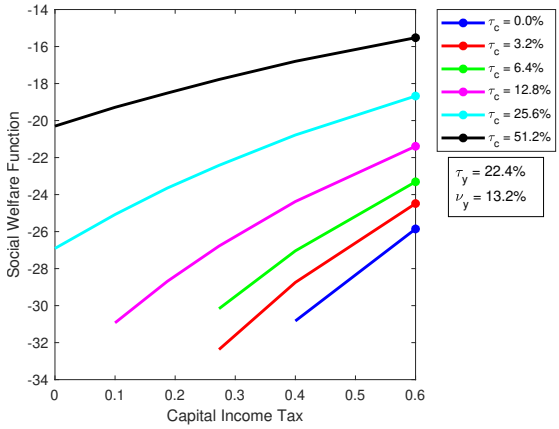
(a) By τ_y



(b) By ν_y



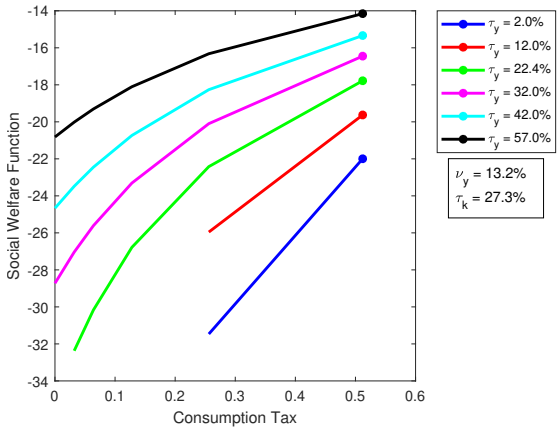
(c) By τ_c



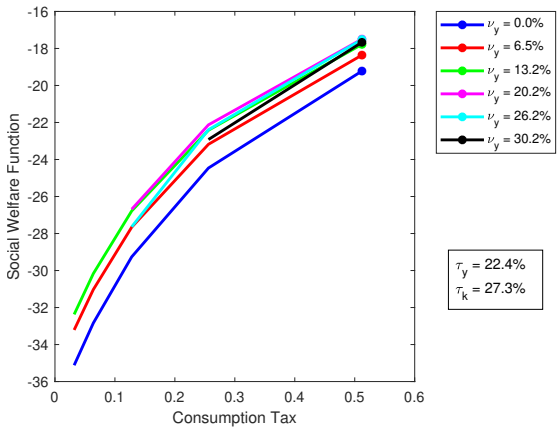
Notes: The figure shows the values for the utilitarian social welfare function along the capital income tax grid for each of the other three tax instruments. Each panel shows the function for different levels for a given tax rate with the other two fixed at their calibrated benchmark values. The top panel depicts it for the labor tax, the center panel for the progressivity parameter, and the bottom panel for the consumption tax. Each line has a highlighted dot that indicates the welfare-maximizing level of capital income at the given tax rate.

Figure 14: Optimal τ_c

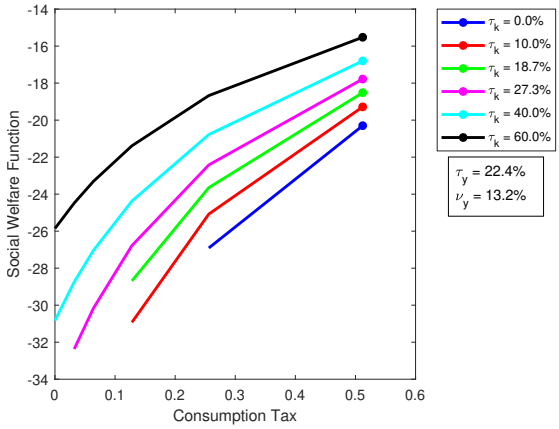
(a) By τ_y



(b) By ν_y



(c) By τ_k

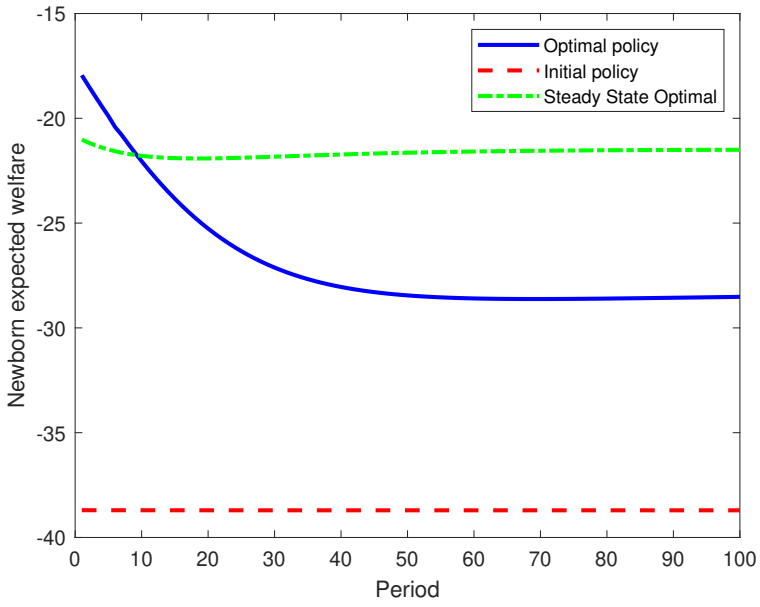


Notes: The figure shows the values for the utilitarian social welfare function along the consumption tax grid for each of the other three tax instruments. Each panel shows the function for different levels for a given tax rate with the other two fixed at their calibrated benchmark values. The top panel depicts it for the labor tax, the center panel for the progressivity parameter, and the bottom panel for the capital income tax. Each line has a highlighted dot that indicates the welfare-maximizing level of consumption tax at the given tax rate.

D.2 Expected Welfare Across Cohorts

Figure 16 plots the expected newborn welfare across cohorts and three different policies.

Figure 16: Expected newborn welfare across cohorts and policies.



Notes: The figure plots the expected welfare of every cohort of young households born along the first 100 periods of transition. Specifically, for each cohort, the figure shows expected newborn welfare for three different policy reforms. "Optimal policy" refers to the reform that maximizes average utility of initial household (i.e., "Baseline" from Table 5). "Steady state optimal" refers to the reform that generates the highest average welfare across steady states (i.e., "Steady state only" from Table 5). Finally, "Initial policy" is the benchmark economy policy. The x-axis displays the number of periods after the initial period in which the cohort arrives. Zero corresponds to the initial young households.