

MANCHESTER
1824

The University of Manchester

Economics

Discussion Paper

Series

EDP-2210

Taxing Consumption in Unequal Economies

Patrick Macnamara, Myroslav Pidkuyko, Raffaele Rossi

November 2022

Economics

School of Social Sciences

The University of Manchester

Manchester M13 9PL

Taxing Consumption in Unequal Economies^{*}

Patrick Macnamara[†]

U. of Manchester

Myroslav Pidkuyko[‡]

Banco de España

Raffaele Rossi[§]

U. of Manchester

November 2022

Abstract

This paper shows that linear consumption taxes are a powerful tool to implement efficient redistribution. We derive this result in an estimated life-cycle economy with labor and capital income risk that reproduces the distribution of income and wealth in the United States. Optimal policy calls for raising all fiscal revenues from consumption, and providing social insurance via a highly progressive wage tax schedule. Capital income and wealth should not be taxed. This policy reduces inequality and increases productivity, and brings large welfare gains both relative to the status-quo and to the case where consumption is not taxed. More than two-thirds of these gains are due to redistribution. Considering transitional dynamics, we show that our reform also generates large welfare gains in the short run.

JEL Classification: E62, H21, H24.

Keywords: optimal policy, inequality, consumption taxation, life-cycle, entrepreneurs.

^{*}We thank Arpad Abraham, Michael Dobrew, Pawel Doligalski, Markus Poschke, Hakki Yazici and the seminar participants at the University of Barcelona, Bristol, Cardiff, Carlos III in Madrid, CEMFI, Manchester, National University of Singapore, Rochester, Bank of Holland, Bank of Spain, CEF Dallas, MMF Warsaw for useful comments. The views expressed in this paper are those of the authors and do not necessarily represent the views of the Bank of Spain and the Eurosystem.

[†]Email: patrick.macnamara@manchester.ac.uk

[‡]Email: myroslav.pidkuyko@bde.es

[§]Email: raffaele.rossi@manchester.ac.uk

1 Introduction

What is the best way to provide redistribution in unequal societies? Increased inequality in the United States makes our question essential in public policy, and our contribution is to show that linear consumption taxes are a powerful tool to provide redistribution without compromising on the size of the economy. Specifically, optimal policy calls for raising all fiscal revenues from consumption, and providing social insurance via a highly progressive labor income tax, while capital income and wealth should not be taxed. Using cross-section data for the US and a quantitative life-cycle model with uninsurable labor and capital income risk, we show that our tax reform proposal increases productivity and reduces inequality and brings about large welfare gains in the sense of [Lucas \(1987\)](#) and [Conesa et al. \(2009\)](#), with around two-thirds of these gains coming from redistribution.

Our life-cycle economy is populated by households who face uninsurable idiosyncratic earning risk, as in standard Aiyagari-Huggett economies. Moreover, the most productive agents can obtain higher-than-average returns on their wealth by choosing to be entrepreneurs and run private businesses (e.g., [Cagetti and De Nardi, 2006](#)). Financial constraints and idiosyncratic fluctuations in entrepreneurial productivity make returns to saving stochastic and heterogeneous, as in [Guvenen et al. \(2022b\)](#). This latter feature is consistent with recent empirical evidence showing a substantial degree of return dispersion (e.g., [Fagereng et al., 2020](#)). It is also a powerful modeling tool that enables us to replicate key features of inequality that have proved challenging to generate through other mechanisms (e.g., [Benhabib and Bisin, 2018](#); [Benhabib et al., 2019](#)). For the aim of this paper, we show that the combination of idiosyncratic labor income risk and financially constrained entrepreneurs makes taxing consumption particularly appealing.

Most of the optimal taxation literature considers consumption taxation alongside income taxes in representative-agent models (e.g., [Coleman, 2000](#); [Correia et al., 2013](#); [Laczó and Rossi, 2020](#)), analyzes consumption taxes in heterogeneous-agent models as an alternative to income taxes (e.g., [Nishiyama and Smetters, 2005](#)), or restricts the tax system to be linear (e.g., [Correia, 2010](#)). Instead, we analyze the welfare benefit of linear consumption taxes in combination with progressive labor income taxes, as well as taxes on capital income and wealth. This enables us to derive our policy conclusions about consumption taxation by considering a wide set of tax instruments commonly in place in modern economies.

In order to best highlight our contribution about consumption taxation, we follow the literature (e.g., [Coleman, 2000](#), [Correia, 2010](#) and subsequent contributions) by restricting aggregate revenues for each tax to be non-negative. This also allows us to concentrate on realistic policies, as aggregate subsidies in labor or capital income are not observed empiri-

cally. Nevertheless, we allow taxes to be negative at the individual level, which is consistent with the US fiscal system (e.g., [Altig et al., 2020](#); [Auerbach et al., 2022](#)).

We find that the optimal policy calls for a consumption tax of around 30% and a zero average tax on labor income. While the labor income tax raises zero revenue in aggregate, the optimal policy calls for a substantial increase in its progressivity (relative to the status-quo), increasing subsidies for the poor and raising marginal tax rates for the rich. In the optimal policy, the richest 1% of wage earners would face a marginal tax rate of 70% and an average tax rate of around 55%, which are substantially higher than their empirical counterparts, 53% and 44%, respectively. Moreover, the ratio between marginal tax wedges of the richest 1% to the bottom 50% (a commonly used measure of progressivity, see [Holter et al., 2019](#)) increases by 45% relative to the status-quo. Meanwhile, wealth and capital income should not be taxed at all. Comparing across stationary equilibria, the welfare gains of the optimal policy are large and amount to 18% of per-capita per-year permanent consumption increase (CEV). Of these gains, around 13% is due to re-distributional effects and slightly over 4% to level effects. Therefore, our proposed policy improves the trade-off faced by governments, in the sense that policymakers can boost redistribution without compromising—actually increasing—the size of the economy.

In order to understand our results, it is instructive to break down the optimal policy exercise into three partial reforms. First, taxing consumption rather than capital income implicitly reallocates resources towards the most productive entrepreneurs who can now expand their businesses with higher-than-average returns. This policy endogenously increases capital intensity, allocative efficiency, the tax base, output and real wages. Due to general equilibrium price effects, virtually all households experience large gains from this reform (around 7% in CEV). To the best of our knowledge, this is the first study highlighting the efficiency gains of taxing consumption on the cross-sectional allocation of capital, in an environment with financially constrained firms.

Second, by taxing consumption rather than wages (but keeping progressivity of the labor income tax schedule fixed at the status-quo), the planner delays the timing of tax extraction, which is welfare improving as long as the interest rate is larger than the population growth rate (e.g., [Summers, 1981](#)). A previously unexplored feature of this policy is that it reallocates wealth towards younger cohorts, thus increasing their self-insurance capabilities in presence of incomplete markets. Again, we show that virtually all households benefit from this and the combined gains of the first two reforms are large (over 12% in CEV). As a result, while the majority of the welfare gains in our optimal policy are due to redistribution, we nevertheless show that two-thirds of the welfare gains could be achieved without increasing the progressivity of the current tax system. Therefore, an important message of our analysis

is that regardless of the desired level of redistribution in the economy, substantial gains can be made through increased use of consumption taxes.

Finally, relative to the status-quo, taxing consumption increases substantially the fiscal space of the government. Strikingly, we find that this on its own creates a scope for increasing social insurance and redistribution via a higher progressive wage tax schedule. While it is well known that incomplete markets models call for social insurance and redistribution in earnings (e.g., [Conesa et al., 2009](#)), what is original here is to show that the optimal degree of progressivity depends importantly on the set of tax instruments available and not exclusively on the underlying income process. This last step further increases the overall welfare gains to around 18% in CEV. Given the large increase in progressivity, around two-thirds of the benefits are due to redistribution effects, with individuals at the top of the distribution experiencing a welfare loss from this reform.

In order to fully isolate the merit of consumption taxation, we also study optimal policy in a scenario where the government cannot tax consumption optimally. With consumption taxes restricted to be equal to the benchmark value of 7.5%, optimal policy calls for taxing wealth rather than capital income, as found by [Guvenen et al. \(2022a,b\)](#). We show that relative to consumption taxation, wealth taxes are a powerful tool to redistribute capital to high-productivity entrepreneurs, but are more detrimental for aggregate wealth accumulation, as taxing wealth erodes the principle. This implies a lower general equilibrium raise in wages (as capital intensity increases less), leading to lower aggregate welfare gains relative to our benchmark exercise.

This result contributes to the current debate about the desirability of a wealth tax in models with financially constrained entrepreneurs and heterogeneous returns (e.g., [Guvenen et al., 2022a](#); [Boar and Midrigan, 2022](#)). Our main point on the matter is that if the government can tax consumption optimally, it decides not to tax wealth (nor capital income). Interestingly, the relative advantages of taxing consumption presented here are above and beyond standard practical issues of implementability of wealth and capital taxation (not modeled here). A brief shortlist includes: increasing capital flight, hidden assets, taxing unrealized capital gains, indivisibility of wealth, distinction between book and market value of wealth, et cetera. Consumption taxation avoids all these issues and as such also has a great appeal for implementation purposes.

We also show that when the government cannot tax consumption optimally, then it relies on both labor and wealth taxes to raise revenue. As a result, the overall fiscal space is much reduced. This decreases the optimal progressivity of the wage tax schedule. In our benchmark exercise with consumption taxation, progressivity is about 23% higher than in this alternative scenario. In summary, in this scenario, the government provides lower

benefits from capital reallocation and only limited social insurance and redistribution in earnings, thus leading to lower benefits in both level and redistribution. We show that a government will lose around 50 percent of the overall welfare gains (or 9% in CEV) by not taxing consumption optimally.

We also discuss the welfare properties of taxing consumption along the transition path between the status quo and the long-run stationary equilibrium. Transitional dynamics are crucial in life-cycle models with a rich demographic structure, as the effects of a policy change on future generations might be radically different than on those alive at the time of the reform (e.g., [Auerbach and Kotlikoff, 1987](#)). In order to precisely evaluate the intergenerational effects of our reform, we follow the literature (e.g., [Auerbach et al., 1983](#); [Altig et al., 2001](#); [Nishiyama and Smetters, 2005](#)), and analyze whether adopting our proposed policy generates welfare gains even after all individuals on impact have been compensated for their potential losses. Thus, we ask whether our tax reform can improve ex-ante welfare in a Pareto sense. Even after compensating those who lose on impact, financed by higher taxes in the future, newborns along the transition experience positive ex-ante welfare gains. Thus, we conclude that our policy is beneficial also in the short run.

The results from this exercise imply that the benefits from taxing consumption are robust to transitional concerns. Therefore we conclude that shifting the burden of taxation towards consumption is appealing from an efficient redistribution point of view, brings large welfare gains, and entails fiscal reforms that are relatively straightforward to implement.

The remainder of the paper is the following. [Section 2](#) presents the model economy. [Section 3](#) describes the estimation procedure. [Section 4](#) presents the optimal policy results. Finally, [Section 5](#) concludes.

2 Model

We present an incomplete-markets life-cycle model consisting of households, firms and a government who interact in competitive good and factor markets.

2.1 Households

The economy is populated by a continuum of households, who differ by age, labor productivity and entrepreneurial ability. Each period, a mass of new households is born, where the rate of population growth is exogenous and assumed to be n . At birth, households learn their type $i \in \{1, \dots, I\}$, which will index its overall level of labor earnings. During their life, households choose consumption, savings, and labor supply and whether or not to engage

in entrepreneurial activity. Households also pay progressive taxes on total income and flat social security taxes on labor earnings (up to a cap). After retirement at age R , households receive social security benefits from the government.

Households also face a risk of early death. We denote by s_j the probability of surviving to age j , conditional on surviving to age $j - 1$, where $s_1 = 1$ and $s_{J+1} = 0$. The demographic patterns are stable, so that age- j agents make up a constant fraction μ_j of the total population. Accidental bequests of type- i are redistributed to all living type- i consumers as a lump-sum transfer, $T_{b,i}$. However we experiment with different assumptions about bequests, including explicit bequests with inter-generational links in labor and entrepreneurial abilities, and results do not change for all practical purposes.

Preferences All agents have identical preferences for consumption c_j and hours worked h_j over their lifetime:

$$E \left\{ \sum_{j=1}^J \beta^{j-1} \left(\prod_{k=1}^j s_k \right) u(c_j, h_j) \right\}, \quad (1)$$

where $\prod_{k=1}^j s_k$ is the unconditional probability an age-1 agent will survive to age j . As it is standard in the literature (e.g., [Conesa et al., 2009](#)), we assume that the period utility is of the form

$$u(c, h) = \frac{[c^\gamma(1-h)^{1-\gamma}]^{1-\sigma}}{1-\sigma},$$

where γ is the consumption utility share and σ controls the household's risk aversion.¹

Labor Earnings Risk In each period before retirement, agents receive labor earnings equal to weh , where w is the real wage rate, e is the household's labor ability and h is hours worked. When households reach age R , they retire so that hours worked and total labor earnings become zero for ages $j \geq R$.

We assume ex-ante and ex-post heterogeneity in labor abilities as in, inter alia, [Kaplan and Violante \(2014\)](#) and [Güvönen et al. \(2022b\)](#). During its working life ($j < R$), a household's labor ability $e_{i,j}(z_h)$ is given by

$$\log e_{i,j}(z_h) = \alpha_0 + \alpha_1 j + \alpha_2 j^2 + \alpha_3 j^3 + \alpha_4 j^4 + \bar{e}_i + \log z_h \quad (2)$$

A household's labor productivity depends on three factors. First, labor ability explicitly depends on a deterministic age profile common across all agents, that we model as a fourth-order polynomial in age j . Second, labor ability depends on an innate, household-specific

¹Given the assumption of a Cobb-Douglas utility function, the coefficient of relative risk aversion in consumption is $-cu_{cc}/u_c = 1 - \gamma(1 - \sigma)$.

fixed effect, \bar{e}_i . At birth, the household learns her type $i \in \{1, \dots, I\}$ which indexes its overall level of labor ability throughout her entire working life. We assume that \bar{e}_i is drawn from a discretized normal distribution with mean zero and variance σ_e^2 , where π_i the probability a household will become type i .² Third, labor ability is also affected by a stochastic idiosyncratic component, z_h , which follows a random walk

$$\log z'_h = \log z_h + \varepsilon_h, \quad \varepsilon_h \sim N(0, \sigma_{\varepsilon h}^2), \quad (3)$$

where the initial $\log z_h$ is set to zero.

Asset Return Risk Through Entrepreneurship All households can choose whether to be an entrepreneur. They have access to a “backyard technology” where they use k units of capital to produce q units of an intermediate capital service according to a linear technology:

$$q = x_j(z_r)k \quad (4)$$

where $x_j(z_r)$ can be interpreted as the household’s entrepreneurial productivity.³ A household’s entrepreneurial productivity, $x_j(z_r)$, is given by

$$\log x_j(z_r) = \kappa_1 \left(j - \frac{1}{J} \sum_{j'=1}^J j' \right) + \kappa_2 \left(j^2 - \frac{1}{J} \sum_{j'=1}^J (j')^2 \right) + \log z_r.$$

Entrepreneurial productivity depends on two factors. First, it depends on a deterministic common age component, that we model as a second-order polynomial in age j .⁴ Second, it is affected by an idiosyncratic shock, z_r , which follows an AR(1) process:

$$\log z'_r = \rho_r \log z_r + \varepsilon_r, \quad \varepsilon_r \sim N(0, \sigma_{\varepsilon r}^2) \quad (5)$$

where the initial shock is drawn from the distribution $N(0, \sigma_{\varepsilon r}^2 / (1 - \rho_r^2))$. The intermediate capital service is then sold at the price p in a perfectly competitive market to the final goods producer, where it is used (along with labor) to produce the uniform final good Y (see Section 2.2 below).

All households lend on the bond market their whole wealth at the riskless rate r . Those who also choose to be entrepreneurs borrow at rate r on the same market and use their own

²Formally, to determine $\{\pi_i\}_{i=1}^I$, we construct a discrete approximation of $N(0, \sigma_e^2)$ using Tauchen (1986).

³Our setup is isomorphic to an alternative setup where entrepreneurs produce the final output good and hire labor.

⁴The level of the age component is chosen so that its average is zero.

backyard technology to produce the intermediate capital service q . Entrepreneurs must also decide how much capital k to invest in their backyard technology. They are subject to a collateral constraint, i.e., $k \leq \lambda a$, where $\lambda \geq 1$ is exogenous and controls the leverage level, while a is the individual entrepreneur's wealth (e.g., see [Moll, 2014](#), [Boar and Midrigan, 2019](#) and [Guvenen et al., 2022b](#)). Entrepreneurs then maximize the following profit function,

$$\pi(a, x) = \max_{0 \leq k \leq \lambda a} \{pxk - (r + \delta)k\}, \quad (6)$$

where p is the price of the capital service, $r + \delta$ is the rental rate of capital, with δ representing the depreciation rate. The associated optimal capital demand is

$$k(a, x) = \begin{cases} \lambda a & \text{if } x \geq (r + \delta)/p \\ 0 & \text{if } x < (r + \delta)/p \end{cases} \quad (7)$$

Therefore, there exists an endogenous productivity threshold,

$$\bar{x} = (r + \delta)/p, \quad (8)$$

such that only households that are sufficiently productive will choose to be entrepreneurs, while the others will simply engage in lending activities. This feature derives from our assumption of constant returns to scale and it allows the model to match the entrepreneurship rate observed in the data. Our framework also avoids the negative relationship between wealth and returns, which is counterfactual (e.g., see [Bach et al., 2020](#)).

To summarize, all households earn the interest rate r by lending their wealth on the bond market. Those households with sufficiently high entrepreneurial ability also choose to run a business, whereby they borrow at rate r , produce the intermediate good q and earn $\pi(a, x)$. Substituting the solution for $\pi(a, x)$, the household's total return on its wealth, $r_a = r + \pi/a$, is given by

$$r_j^a(z_r) = r + \lambda \max(px_j(z_r) - (r + \delta), 0). \quad (9)$$

Therefore, there will be persistent idiosyncratic variation in returns across households, which is a crucial ingredient for the model's ability to match the fat tail of wealth and taxable income (e.g., see [Benhabib et al., 2011](#), [Benhabib et al., 2019](#) and [Guvenen et al., 2022b](#)). Furthermore, despite no explicit link between wealth and returns, high-wealth households will, on average, earn higher returns, consistent with the empirical evidence (e.g., see [Bach et al., 2020](#) and [Fagereng et al., 2020](#)).

2.2 Final Production Firm

The final good is produced according to a Cobb-Douglas production function:

$$Y = F(Q, L) = Q^\alpha L^{1-\alpha}$$

where L is aggregate labor and Q is the aggregate of the intermediate capital good produced by entrepreneurs.

It is straightforward to derive the following aggregate relationship:

$$Y = AK^\alpha L^{1-\alpha}$$

where K is aggregate capital and A is aggregate TFP. Aggregate TFP is $A = (Q/K)^\alpha$, where Q/K is the average productivity of entrepreneurs. Therefore, aggregate productivity depends crucially on the allocation of capital across entrepreneurs.

The market for the intermediate good and the market for labor are both perfectly competitive. Therefore, the representative firm takes as given the prices (w, p) and chooses Q and L to maximize profits, $\Pi = Q^\alpha L^{1-\alpha} - pQ - wL$.

2.3 Government

The government taxes income in order to finance a fixed and exogenous level of government spending, G , which provides agents no utility. The government operates a balanced budget and does not use debt, implying that G is just equal to aggregate tax revenue. The government also runs a social security system with a dedicated budget.

2.3.1 Taxes

The government obtains revenue from several potential sources: (1) a progressive labor income tax, (2) a flat capital income tax, (3) a flat consumption tax, and (4) a flat wealth tax.

Labor Income Tax Labor and capital income are separately taxable. Households can deduct part of the social security contribution from their labor income, up to an upper limit \bar{y} . The household's taxable labor income

$$y_l = we_{i,j}(z_h)h - \frac{\tau_{ss}}{2} \min (we_{i,j}(z_h)h, \bar{y})$$

We specify a progressive income tax schedule, following [Heathcote et al. \(2017\)](#), where the household's total labor income tax is

$$T_l(y_l) = y_l - \lambda_l y_l^{1-\tau_l}. \quad (10)$$

The progressivity of the income tax schedule is controlled by τ_l , and the level of income taxes is determined by λ_l . As such, the functional form of the tax function implied by (10) permits a precise measure of tax progressivity that is not confounded by the level of tax rates. We also present a second measure of progressivity that is commonly used in the literature (e.g., [Güvenen et al., 2014](#); [Holter et al., 2019](#)) and does not rely on a specific functional form of the tax function. This reads as

$$PW = 1 - \frac{1 - T'_l(y_2)}{1 - T'_l(y_1)}, \quad (11)$$

where $T'_l(y_1)$ is the marginal tax rate paid by an household with labor income y_1 . The variable PW takes values between 0 and 1, as long as the tax schedule is weakly progressive. Conveniently, given the tax function in (10), we can rewrite (11) as

$$PW = 1 - \frac{1 - T'_l(y_2)}{1 - T'_l(y_1)} = 1 - \left(\frac{y_1}{y_2}\right)^{\tau_l},$$

for any arbitrary labor income levels $y_2 \geq y_1$. For practical purposes in our exercise y_1 is the median wage income and y_2 is the marginal income necessary to be in top 1 percent of the labor income distribution.

Capital Income Tax Household's taxable capital income is

$$y_k = r_j^a(z_r)a \quad (12)$$

We assume a flat tax on capital income, at the rate τ_k .

Consumption Tax We assume a flat consumption tax, at rate τ_c .

Wealth Tax We assume a flat wealth tax, at rate τ_a . Given wealth a , the household pays a tax $\tau_a a$. In our benchmark economy, we set $\tau_a = 0$, but in our optimal policy analysis, we explicitly allow for a wealth tax.

2.3.2 Social Security Scheme

The government runs a pay-as-you-go social security scheme. Taxpayers pay a social security tax only out of their labor income (at the flat tax rate τ_{ss}), up to an upper bound \bar{y} . Once an individual has reached retirement age ($j \geq R$), the government pays out a social security benefit b_i :

$$b_i = \chi \Phi(\min\{wL_i, \bar{y}\})$$

where χ is a parameter which ensures the social security budget constraint is satisfied.⁵ $\Phi(\cdot)$ is a progressive function of the household's average labor income for his type (below the cap \bar{y}). The function is modeled to be consistent with the US Social Security benefit schedule:

$$\Phi(y) = \begin{cases} 0.9y & \text{if } y \leq y_1 \\ 0.9y_1 + 0.32(y - y_1) & \text{if } y_1 < y \leq y_2 \\ 0.9y_1 + 0.32(y_2 - y_1) + 0.15(y - y_2) & \text{if } y_2 < y \leq \bar{y} \end{cases}$$

This function depends on two bend points (y_1, y_2) as it is in the US Social Security .

Social security benefits are financed by a flat tax τ_{ss} on all labor earnings weh below \bar{y} . That is, a household with labor earnings weh will pay a social security tax of $\tau_{ss} \min(weh, \bar{y})$. Given the tax rate τ_{ss} and the cap \bar{y} , we internally set the parameter χ so that aggregate social security tax revenue equals aggregate social security benefits.

2.4 Value Function

Having presented the main features of our model economy, we can now describe the household's problem in recursive form. In each period, the household chooses consumption c , savings a' , and labor supply h given idiosyncratic risk, the sequence of prices and taxes. In retirement, households supply zero hours (i.e., $h = 0$), but they still choose consumption and savings. Let $V_{i,j}(a, z_h, z_r)$ denote the value of a type- i and age- j consumer with assets a and idiosyncratic shocks (z_h, z_r) . We can write the consumer's maximization problem as follows:

$$V_{i,j}(a, z_h, z_r) = \max_{c, h, a'} \{u(c, h) + \beta s_{j+1} E[V_{i,j+1}(a', z'_h, z'_r) | z_h, z_r]\} \quad (13)$$

⁵In our benchmark economy, χ turns out to be 0.985.

subject to

$$\begin{aligned}
(1 + \tau_c)c + a' &= (1 - \tau_a)a + (1 - \tau_k)r_j^a(z_r)a + we_{i,j}(z_h)h \\
&\quad - \tau_{ss} \min (we_{i,j}(z_h)h, \bar{y}) - \mathcal{T}_l(y_l) + T_{b,i} + b_i \mathbb{1}\{j \geq R\} \\
y_l &= we_{i,j}(z_h)h - \frac{\tau_{ss}}{2} \min (we_{i,j}(z_h)h, \bar{y}) \\
a' &\geq 0 \\
0 &\leq h \leq \mathbb{1}\{j < R\}.
\end{aligned}$$

2.5 Equilibrium

We focus on a stationary equilibrium, in which capital, labor, transfers and government consumption are all constant in per-capita terms. See Appendix A for a full definition of the equilibrium.

3 Quantitative Analysis

We describe here our estimation approach, and then evaluate the model’s ability to account for a number of features in the data for the US. We solve and estimate the model assuming the economy is in a steady state. One period corresponds to one year and we convert all nominal values into 2010 dollars. Our numerical strategy is described in Appendix B.

Following a long tradition in structural public finance (e.g., [Gourinchas and Parker, 2002](#); [French and Jones, 2011](#)), we adopt a two-step estimation procedure. This consists of dividing our parameters into two main groups: (i) a group of parameters that is externally set, either according to previous literature, via direct observation or through estimation; and (ii) a group of parameters that is internally set, estimated using a Simulated Method of Moments (SMM) estimator, in order to match relevant distributional moments in the Survey of Consumer Finances (SCF) for 2019 and other standard macroeconomic moments from national accounts. This is by now a standard methodology in applied structural works in macroeconomics (e.g., [Heathcote et al., 2014](#); [Benhabib et al., 2019](#)).

3.1 Externally Set Parameters

All parameters externally set are reported in Table 1.

Externally Fixed Parameters We set the parameter governing the agents’ risk aversion, σ , to 4, which is a standard value for life-cycle economies, (e.g., [Conesa et al., 2009](#); [Güvönen](#)

et al., 2022b). We fix our Cobb-Douglas parameter in order to recollect the income shares from NIPA (i.e., $\alpha = 0.36$).

Then, we fix J , the maximum age in the model, to 85 and R , the retirement age, to 45. Assuming that age 1 in the model corresponds to age 21 in the real life, these choices for (J, R) correspond to ages 105 and 65 in real life/years. We set the population growth rate n to 0.7%, to be consistent with the US population growth rate in the World Bank’s World Development Indicators. We obtained estimates of the survival probabilities s_j from the United States Mortality Database (see Appendix C.1 for details).

Next, we set the two level of income, y_1 and y_2 of the Social Security function $\Phi(y)$ according to the corresponding bend points values obtained from the Table on the Social Security Website.⁶ We set the status-quo consumption tax rate, $\tau_c = 0.075$ and capital income tax rate, $\tau_k = 0.25$. This follows Guvenen et al. (2022b) who base their calibration on McDaniel (2007). And finally, as the US has no wealth tax, we set $\tau_a = 0$ in our benchmark economy.

Externally Estimated Parameters Here we focus on a set of parameters that we directly estimate outside the model (see panel B, Table 1). We begin by estimating the progressivity of the non-linear tax function on labor earnings (τ_l).⁷ Using our SCF data, we construct a measure of labor income and then calculate federal income tax liabilities using NBER’s TAXSIM program, but assume the household has no capital income. See Appendix C.2 for details.

Next we estimate the labor productivity process. First, we compute our quartic age-earnings profile from SCF. This allows us to directly estimate the α_j parameters, as reported in Table 1. Then, we estimate the other two parameters of interest for the labor productivity process internally within our SMM routine.

3.2 Internally Estimated Parameters

We use SMM to estimate the remaining twelve parameters, $(\gamma, \alpha_0, \sigma_e, \sigma_{\varepsilon h}, \lambda_l, \kappa_1, \kappa_2, \rho_r, \sigma_{\varepsilon r}, \beta, \lambda, \delta)$. Briefly, this estimator consists of choosing the structural parameters such that the moments computed from real data are as close as possible to those computed from data simulated from our model (e.g., Gourinchas and Parker, 2002; Cagetti and De Nardi, 2006; Heathcote et al., 2014; Benhabib et al., 2019). In particular, indicating the vector of parameters to be estimated by Θ , the SMM estimator solves the following minimum distance

⁶<https://www.ssa.gov/oact/cola/bendpoints.html>.

⁷We will estimate the level of the tax function, λ_l internally. To identify this parameter, we target the average labor income tax rate.

Table 1 – Externally Set Parameters

Parameters	Notation	Value	Std. Err.	Source
<u>A: Fixed Parameters</u>				
Risk Aversion	σ	4		Typical in lit.
Capital Share	α	0.36		Typical in lit.
Maximum Age	J	85		Corresp. to age 105
Retirement Age	R	45		Corresp. to age 65
Survival Prob.	s_j	Appendix C.1		USMD 2018
Pop. Growth	n	0.007		World Bank
Soc. Sec. Tax	τ_{ss}	0.124		IRS
Soc. Sec. Bend Pt. 1	y_1	9.33		SSA
Soc. Sec. Bend Pt. 2	y_2	56.23		SSA
Soc. Sec. Cap	\bar{y}	107.7		IRS
Soc. Sec. Benefit	χ	0.985		Balanced budget
Cons. Tax Rate	τ_c	0.075		Güvönen et al. (2022b)
Cap. Income Tax Rate	τ_k	0.25		Güvönen et al. (2022b)
Wealth Tax Rate	τ_a	0		No Wealth Tax in US
<u>B: Estimated Parameters</u>				
Labor Tax, Prog.	τ_l	0.20	(0.0014)	TAXSIM
Ability Coef. 1	α_1	0.147	(0.013)	SCF 2019
Ability Coef. 2 ($\times 10^3$)	α_2	-7.25	(1.13)	SCF 2019
Ability Coef. 3 ($\times 10^4$)	α_3	1.66	(0.37)	SCF 2019
Ability Coef. 4 ($\times 10^6$)	α_4	-1.43	(0.41)	SCF 2019

Note: This table reports the externally set parameters. USMD stands for the United States Mortality Database. Standard errors are reported in parentheses.

problem:

$$\hat{\Theta} = \arg \min_{\Theta} \left(\hat{M} - \hat{m}(\Theta) \right)' W \left(\hat{M} - \hat{m}(\Theta) \right), \quad (14)$$

where \hat{M} denoted the targeted cross-sectional moments from the 2019 SCF as well as macroeconomic moments from standard National Income Accounts and [Jordà et al. \(2019\)](#). The matrix $\hat{m}(\Theta)$ represents the moments implied by the model for a given set of parameters Θ , and W is a weighting matrix.⁸

The estimated parameters are reported in the top panel of [Table 2](#), while the moments are reported in bottom panel of the same table. We now give insights about our identification

⁸We freely picked the weighting matrix W . In particular, we assumed the off-diagonal elements are all zero. For the diagonal elements, we assume $W_{ii} = 1/\hat{M}_i^2$, where \hat{M}_i is data moment i . This approach is common in the literature, in light of the Monte Carlo results presented by [Altonji and Segal \(1996\)](#), who argue that in standard applications there is a non-negligible small sample bias when using the optimal weighting matrix.

strategy by describing how the structural parameters are useful for capturing various aspects of the data moments. While this is not a precise mapping, as the estimation captures a large number of spillovers and general equilibrium price effects, we believe it is still important to generate useful intuition.

First, the preference parameter controlling the utility weight of consumption, γ , is useful for matching labor supply moments, such as average hours. This is because this parameter indirectly controls labor supply, via the importance of leisure. Next, the constant term in the earning ability age profile, α_0 , act as a scaling factor and is useful for capturing average earnings. The parameters σ_e and $\sigma_{\varepsilon h}$ are the standard deviations of the innate labor abilities and idiosyncratic ability, respectively. These two parameters are useful to identify the remaining moments regarding the earning distribution, such as the earnings gini, the slope of the earnings gini by age and the earnings share at the top of the distribution. The level parameter for the labor tax function, λ_l , is useful for matching the average tax rate on labor income.

Next, the parameters on the age profile of the entrepreneurial ability (κ_1, κ_2) as well as the two parameters governing the AR(1) process of the idiosyncratic component of ($\rho_r, \sigma_{\varepsilon r}$) of the same abilities, are crucial both for matching the right tail of the wealth distribution, as well as for matching various important moments of entrepreneurial data (such as age and average span of entrepreneurial activity).

Finally, we report standard macroeconomic parameters. The discount factor, β , is useful for matching in the aggregate capital-income-ratio. The parameter controlling the maximum entrepreneurial leverage ratio, λ , helps in matching the risk-free interest rate rate, while the capital depreciation rate, δ , enables the model to capture the economy-wide investment-to-output ratio.

Furthermore, the model captures extremely well aggregate macroeconomic data and crucially, distributional data. This is pivotal for the aim of studying the aggregate and distributional properties of tax reforms in unequal economies. In particular, our model matches the wealth gini and the wealth shares of the wealthiest top 1, 5 and 20 percent, respectively. Similarly, our model matches the right tail in the distribution of earnings. It is interesting the ability of the model of capturing various entrepreneurial characteristics. This is important for the aim of the paper, given that our microfoundation of the right tail of the distribution is based on entrepreneurial activity.

All parameters are statistically different from zero and precisely estimated. This finding is not obvious and shows a tight link between the targeted moments and structural parameters. As parameter identification in SMM requires choosing moments whose predicted values are sensitive to the model's underlying parameters, the results presented here indicate that we

Table 2 – Estimated Parameters and Targeted Moments

Parameters	Notation	Value	Std. Err.
Utility Cons. Weight	γ	0.374	(0.004)
Labor Ability Constant	α_0	2.851	(0.099)
Std. Dev. of Perm. Lab. Ability	σ_e	0.523	(0.021)
Std. Dev. of Idios. Lab. Shock	$\sigma_{\epsilon h}$	0.215	(0.005)
Labor Tax, Level	λ_l	2.220	(0.008)
Ent. Ability, Coef 1 (x100)	κ_1	2.451	(0.214)
Ent. Ability, Coef 2 (x10000)	κ_2	-3.156	(0.122)
Return Persistence	ρ_r	0.988	(0.001)
Return Shock	$\sigma_{\epsilon r}$	0.117	(0.010)
Discount Factor	β	0.995	(0.004)
Coll. Constraint	λ	2.518	(0.120)
Depreciation Rate	δ	0.045	(0.002)

Moments	Model	Data
<i>Cross-Sectional Moments</i>		
Wealth Gini	0.817	0.853
Wealth Share, Top 1%	0.372	0.373
Wealth Share, Top 5%	0.585	0.650
Wealth Share, Top 20%	0.836	0.874
Earnings Gini	0.635	0.649
Earnings Gini-Age Slope	0.738	0.747
Earnings Share, Top 1%	0.120	0.120
Earnings Share, Top 5%	0.293	0.279
Earnings Share, Top 20%	0.568	0.575
Average Earnings	58.08	57.84
Average Labor Tax	0.135	0.135
Average Hours (working age)	0.310	0.313
Entrepreneurship Rate	0.088	0.092
Ent. Rate Age, Coef 1 (x100)	0.770	0.725
Ent. Rate Age, Coef 2 (x10000)	-1.001	-1.006
Avg. Numb. of Years an Ent.	14.44	14.67
Avg. Age Became an Ent. (Real Life Age)	39.44	39.33
<i>Macroeconomic Moments</i>		
Capital-to-output Ratio	2.72	2.95
Investment-to-output Ratio	0.21	0.22
Borrowing Rate	0.019	0.019

Note: The top panel reports the estimated parameters with standard errors, while the bottom panel reports the moments in the model and the data. The model parameters are estimated via Simulated Method of Moments (SMM) using moments from the 2019 Survey of Consumer Finances (SCF).

picked the right targets.

The estimated parameters have values that are, broadly speaking, consistent with those found in the literature. This is the case for the discount factor β , the utility parameter γ , the collateral constraint λ and the depreciation rate δ . Similarly, our parameters governing earning dynamics are broadly consistent with the existing literature estimating similar processes, e.g., [Kaplan and Violante \(2014\)](#) and [Guvenen et al. \(2021\)](#). We have no good prior for the parameters governing the return profiles, although recent quantitative studies point to substantially persistent and moderately variable processes (e.g., [Xavier, 2020](#) and [Guvenen et al., 2022b](#)). As such, our estimates are consistent with these results.

3.3 Model’s Performance on Untargeted Moments

In order to increase the plausibility of how we discipline our model, we run two external validity exercises on untargeted statistics. First, we show that the return profiles at the top of the wealth distribution generated by our model are consistent with those found in the data. Second, we show that our model matches the capital income shares at the top of the income distribution.

Return Profiles The first aspect of our framework that we want to assess is to what extent the return heterogeneity necessary to capture distributional moments in entrepreneurial moments and wealth distribution is somewhat consistent with the empirical evidence on wealth returns. This is not an easy task as a number of studies have shown various measurement drawbacks for wealth data and their return in SCF (e.g. [Bhandari et al., 2020](#) and [Smith et al., 2021](#)). One could use recent evidence from Scandinavian countries (e.g. [Fagereng et al., 2020](#)). However, it appears that the portfolio composition of wealth (particularly at the top of the distribution) in the US is quite different than in Northern Europe.⁹

To partly overcome these issues, we compute a rough measure of returns by wealth percentile with the reported assets in SCF. Briefly, we estimate the return for each household’s portfolio using outside estimates of the return on individual asset classes and assuming equal returns within each asset class (e.g., [Saez and Zucman, 2016](#)).¹⁰ We use the estimates of the average returns of different asset types between 1990 and 2019, as reported by [Xavier \(2020\)](#). [Table 3](#) presents the returns by wealth percentiles in the model and in the data.

⁹For example, according to [Fagereng et al. \(2020\)](#), in Norway, the share of housing in gross wealth held by the 95-99 percentiles (99-99.9 percentiles) is 0.73 (0.44). Meanwhile, it is only 0.33 (0.25) in the US (as reported in the 2019 SCF).

¹⁰As such, all heterogeneity comes from portfolio composition. This is to some extent unrealistic, but still useful to evaluate the implications of the model for returns at the top of the wealth distribution. For a detailed study on the consequences of heterogeneous returns within asset classes, see [Smith et al. \(2021\)](#).

Table 3 – Return Profiles (Untargeted)

Wealth Percentile	Returns		
	Model	Data	Data w/o H.
[99-100]	0.095	0.074	0.078
[95-99)	0.043	0.066	0.069
[90-95)	0.033	0.059	0.062
[75-90)	0.027	0.053	0.048
[50-75)	0.022	0.049	0.037
[25-50)	0.021	0.040	0.027
[10-25)	0.021	0.021	0.021
[1-10)	0.021	0.028	0.023

Note: This table reports the resulting wealth returns by wealth percentile in the model and the data. For the data, we estimate the average return for each household’s portfolio in the SCF using estimates of the average returns of different asset types between 1990 and 2019, as reported by [Xavier \(2020\)](#).

A few considerations are in order. First, qualitatively, our model captures the positive correlation between wealth and wealth returns, which is a stylized fact reported in most empirical studies ([Fagereng et al., 2020](#), [Bach et al., 2020](#) and [Smith et al., 2021](#)). In our model, high productivity entrepreneurs have a strong incentive to scale up their businesses, and as such they accumulate larger wealth. Second, the consistency of our model-implied returns at the bottom 25 percent of the wealth distribution and, crucially, in the top 1 percent is striking, since these moments were not targeted in the estimation exercise. We also find that the imputed returns from the SCF seem to be higher than those implied by the model in the middle of the wealth distribution. This is mostly due to returns related to housing for which the model abstracts. All in all, the main take home from this exercise is that the return profile generated by our model is generally consistent with the empirical evidence, particularly at the top of the distribution.

Capital Income Shares A second potential concern about our quantitative exercise is whether our model captures the capital income shares along the income distribution found in the data. This is important for at least two reasons. First of all, because recent empirical studies estimate that in the US, the dynamics in income concentration over the past three decades are mainly driven by a boom in capital income at the top (e.g., [Piketty et al., 2018](#)). Second, and perhaps more importantly for the aim of our paper, matching the capital income shares along the distribution is crucial for a sound evaluation of the relative merits of different tax components.

Table 4 reports the capital income shares of the model with the ones in the data. Reassuringly, the model implies shares of capital income in aggregate and along the distribution

Table 4 – Capital Income Share Along Income Distribution for 2016 (Untargeted)

	Model	Data
All Taxpayers	0.28	0.28
Income Top 1%	0.69	0.61
Income Top 5%	0.48	0.46
Income Top 10%	0.40	0.41
Income Bottom 90%	0.09	0.18

Note: Updated data series from [Piketty et al. \(2018\)](#), available on Gabriel Zucman’s website at <https://gabriel-zucman.eu/usdina/>. Aggregate income is GDP minus depreciation of capital. Capital income includes: i) income from equity; ii) net interest payments; iii) income from housing rents; iv) capital component of mixed income; v) property income paid to pensions.

that are broadly consistent with those found in the data. Overall, this confirms the ability of our model to provide a realistic explanation of the distributional features of the US economy and as such to provide a suitable laboratory for sound policy analysis.

4 Results

We now derive our main results about the optimality of consumption taxation. We first compute the optimal policy where the government has access to a rich set of tax instruments: progressive labor income taxation, flat consumption taxes, flat wealth taxes and flat capital income taxes.

For now we will concentrate on steady-state outcomes, so we can abstract from consideration of intergenerational transfers of the tax burden. However, we will analyze transitional dynamics in a dedicated section.¹¹

Welfare Calculation In each step, the merits of optimal policy and various partial tax reforms are analyzed both in terms of equilibrium outcomes of the endogenous variables and from a welfare point of view. On this latter point, we follow a large literature in macroeconomics, and adopt as a welfare metric the famous calculation of [Lucas \(1987\)](#), adapted to a life-cycle environment. In practical terms, the overall welfare change can be written as,

$$\mathcal{W}(c_*, h_*) = \mathcal{W}(c_0(1 + CEV), h_0), \quad (15)$$

where a generic variable x_* identifies its value under the fiscal reform under study, while x_0 represents the variable in the status-quo. The variable CEV is the traditional consumption

¹¹Note that our tax reforms will have an impact on household income and as such, it changes the Social Security scheme. In the remainder of the paper, we keep the Social Security taxes at their benchmark values and adjust the Social Security benefit level so that the Social Security budget remains balanced.

of equivalent variation, measured in per-capita annual consumption terms. This is the main metric of the welfare effects of a given fiscal reform.

Following [Conesa et al. \(2009\)](#), we decompose the overall change of the fiscal reform into the welfare effects due to variations in consumption and leisure, i.e.

$$1 + CEV = (1 + CEV_c)(1 + CEV_h),$$

where

$$\mathcal{W}(c_*, h_0) = \mathcal{W}(c_0(1 + CEV_c), h_0);$$

$$\mathcal{W}(c_*, h_*) = \mathcal{W}(c_*(1 + CEV_h), h_0).$$

We can further decompose the overall welfare effect of a policy reforms on a specific component, say consumption (same thing for leisure, *mutatis mutandis*, see Appendix [D](#)), into its level effect CEV_{cL} and its distribution effect CEV_{cD} ,

$$\mathcal{W}(\hat{c}_0, h_0) = \mathcal{W}(c_0(1 + CEV_{cL}), h_0);$$

$$\mathcal{W}(c_*, h_0) = \mathcal{W}(\hat{c}_0(1 + CEV_{cD}), h_0).$$

where $\hat{c}_0 = (C_*/C_0)c_0$ is the consumption allocation resulting from scaling the allocation c_0 by the change in aggregate consumption C_*/C_0 . We define the level effect of a given reform the as the product of the level effects (same thing for distribution effects, *mutatis mutandis*) of leisure and consumption, i.e.

$$1 + CEV_L = (1 + CEV_{cL})(1 + CEV_{hL}).$$

4.1 Optimal Policy

We now proceed to our main optimal policy experiment. Namely, we solve for the tax policy which maximizes the ex-ante welfare of a newborn in the steady-state of our economy. While we allow the tax authority to use a rich set of tax instruments (progressive labor, flat consumption, flat wealth, flat capital), we constrain the tax policies so that every tax must raise non-negative tax revenue in aggregate. As such we follow the recommendations of [Coleman \(2000\)](#) and [Correia \(2010\)](#), who argue that large aggregate subsidies in labor or capital income would raise serious concerns about the implementability of such policies. As for the labor income tax, while aggregate tax revenue must be non-negative, subsidies at the individual level are allowed. Indeed, while individual subsidies are commonly observed in

modern economies (e.g., [Altig et al., 2020](#)), aggregate subsidies are less so.

In an important contribution, [Correia \(2010\)](#) advocates the equitable nature of consumption taxation in an environment where general equilibrium dynamics are determined by an infinitely-lived representative agent through Gorman aggregation. With this assumption, the model remains more tractable, but it abstracts from the social insurance aspect of taxation and restricts the analysis to linear taxes. Differently, our model admits complex spillovers from the cross-sectional allocation to equilibrium prices and vice versa, so that it violates Gorman aggregation. This allows us to analyze the benefits of consumption taxation alongside linear and non-linear taxes within a rich quantitative life-cycle model of inequality with incomplete markets. This is a crucial aspect of our study because, as shown by the literature (e.g., [Nishiyama and Smetters, 2005](#)), life-cycle considerations and market incompleteness generate a scope for social insurance and redistribution and fundamentally change tax analysis.

Table 5 compares the outcome of this exercise across stationary equilibria, starting from the status-quo described in Section 3. In the optimal policy, only consumption taxes and labor income taxes are used by the government. Moreover, the labor income tax raises zero revenue in aggregate, but the progressivity of wage tax schedule increases relative to the status-quo from 0.20 to 0.35. These numbers imply an increase in the ratio of the tax wedges between the richest top 1% and the poorest 50% of around 45% (i.e., $\frac{PW_{\text{Optimal Policy}}}{PW_{\text{Status-Quo}}} = 1.45$). The richest earners in the top 1% of the distribution would experience an increase in both their average marginal wage tax rate (AMTR), from 50% to around 70%, and their average tax rate (ATR) from 44% to 54%. At the same time, the poorest 50% experience a decrease in the AMTR from 20% to 17.4% and a reduction in ATR from 1% to -24%.

With this policy in place, wealth and its quality-adjusted measure increase sensibly by 16.6% and 27%, respectively. In equilibrium, investment mirrors the behavior of wealth, recording an increase by roughly 19%. Output slightly increases (+0.4%), and, interestingly, hours worked decrease by a substantial amount (12.1%). Moreover, we find an increase in allocative efficiency, that is reflected by the increase in TFP (+3.1%) and in real wages (+14.1%). At the same time, the large boost in wealth decreases the price of capital and the corresponding borrowing rate (i.e., -20.9% and 1 percentage point, respectively).

While this policy increases efficiency as measured by higher wealth and TFP, it is also effective in reducing inequality. The wealth Gini decreases by 2.2%, and most importantly, consumption Gini sees a 10.7% reduction. The optimal policy delivers an 18% welfare gain in consumption equivalent terms, of which 13% is due to increase redistribution (around two-third of the overall gains), while 4% to level gains. Around 89% of newborns are better off with this policy.

Table 5 – Optimal Policy

	Benchmark	Optimal Policy
<i>Policy Rates:</i>		
Consumption Tax	7.5%	30.4%
Avg. Labor Income Tax	13.5%	0.0%
Labor Tax Progressivity (τ_l)	0.20	0.35
Capital Income Tax	25.0%	0.0%
Wealth Tax	0.0%	0.0%
<i>Aggregate Quantities:</i>		
Wealth		16.6
Quality-Adj. Wealth		27.0
Hours		-12.1
Output		0.4
Consumption		-2.9
Investment		18.9
<i>Productivity:</i>		
TFP		3.1
Entrepreneurial Rate (Δ p.p.)		-2.9
<i>Prices:</i>		
Price of Capital		-20.9
Wages		14.1
Borrowing Rate (Δ p.p.)		-1.0
<i>Inequality:</i>		
Wealth Gini		-2.2
Earnings Gini		0.1
Consumption Gini		-10.7
<i>Welfare:</i>		
CE Welfare Gain		18.0
<i>Level</i>		4.0
<i>Redistribution</i>		13.5
Pct. of Newborns Better Off		88.6

Note: In this table, we report the optimal policies (for the ex ante welfare of a newborn), where aggregate tax revenue is constrained to be non-negative for each tax instrument. In the optimal policy, the fiscal authority has access to a flat consumption tax, a progressive labor tax, a flat capital tax and a flat wealth tax.

4.2 Partial Reforms

In order to gain further insights into the sources of these welfare gains, we next decompose the optimal policy into a sequence of partial tax reforms. We outline these reforms below:

1. First, we eliminate the capital income tax.
2. Second, starting from the first reform, we adjust the average labor income tax to zero, but maintain the labor income tax progressivity of the status-quo.
3. Third, from the second reform, we adjust the labor income tax progressivity to the optimal level.

Notice that in the final reform, we move to the optimal policy. In all cases, we use the consumption tax to raise any lost revenue. Nevertheless, in Section 4.3, we will compare our results with an alternative scenario in which the government uses a wealth tax to raise the lost revenue. Table 6 reports the results of these simple reforms.

4.2.1 Partial Reform 1: Eliminate Capital Income Tax

We study here the effects of a tax reform in which the government replaces capital income taxes (i.e., $\tau_k = 0$) with consumption taxes (τ_c). At the same time, we keep taxes on earnings unchanged at the status quo. We compare both aggregate and distributional outcomes in the stationary equilibrium of this alternative fiscal arrangement with those of the benchmark model.

Relative to more complex tax reforms, this exercise is useful for three important reasons. First, it keeps fixed the tax instruments other than consumption taxes (e.g., wage taxes). In this way we can provide a clean scrutiny of how consumption taxes work and how they differ from capital income taxes in isolation of other mechanisms. Second, its relative simplicity also makes it appealing and easy to communicate from a policy perspective. Finally, this policy reforms resembles those proposed and implemented in a number of OECD countries, such as the UK and China. Thus understating its merits has a direct practical value.¹²

Macroeconomic Effects We start by explaining the findings of the second column of Table 6. A few results are worth noting. First, regarding macroeconomic outcomes, aggregate

¹²More generally, in the past five decades, most OECD countries have shifted their tax burden from income to consumption, mainly in the form of a Value Added Tax. Since 1965, in these countries, the share of consumption taxes as a percentage of GDP has more than doubled, from 3.2 percent to 7.1 percent in 2018. At the moment they raise around 33 percent of total tax revenues on average, compared with 11.9 percent in 1965 (e.g., [OECD, 2020](#)).

Table 6 – Partial Reforms

Variable	Benchmark	Eliminate Cap. Inc. Tax	+ Adjust Avg. Lab. Inc. Tax	+ Optimal Policy
<i>Tax Instruments:</i>				
Consumption Tax	7.5%	12.2%	25.8%	30.4%
Avg. Lab. Income Tax	13.5%	15.1%	0%	0%
Labor Tax Progressivity	0.20	0.20	0.20	0.35
Capital Income Tax	25%	0%	0%	0%
Wealth Tax	0%	0%	0%	0%
<i>Aggregate Quantities:</i>				
Wealth		15.7	35.0	16.6
Quality-Adj. Wealth		31.0	43.1	27.0
Hours		-2.1	0.4	-12.1
Output		9.3	14.3	0.4
Consumption		10.6	14.1	-2.9
Investment		21.3	37.5	18.9
<i>Productivity:</i>				
TFP		4.6	2.1	3.1
Entrepreneurial Rate (Δ p.p.)		-4.3	-1.9	-2.9
<i>Prices:</i>				
Price of Capital		-16.6	-20.1	-20.9
Wages		10.8	13.4	14.1
Borrowing Rate (Δ p.p.)		-0.6	-1.0	-1.0
<i>Inequality:</i>				
Wealth Gini		3.6	2.0	-2.2
Earnings Gini (pre-tax)		0.5	0.2	0.1
Consumption Gini		5.7	2.5	-10.7
<i>Welfare:</i>				
CE Welfare Gain		6.9	12.2	18.0
<i>Level</i>		12.4	13.8	4.0
<i>Redistribution</i>		-4.6	-1.4	13.5
Pct. of Newborns Ex-Post Better Off		100	100	88.6

Note: This table reports a how the optimal policy (the last column) can be decomposed into a sequence of revenue-neutral simple tax reforms. All variables (except for the tax instruments) are reported as the percentage or percentage point change relative to the status quo. In the first reform, the capital income tax is eliminated. Starting from the first reform, the second reform reduces the average labor income tax to zero, keeping the progressivity fixed at the status-quo. Starting from the second reform, the final reform increases the labor income tax progressivity to the optimal level. In all cases, the consumption tax is used to recover any lost revenue.

Table 7 – Effect of Simple Reforms on Allocation of Capital

Variable	Eliminate Cap. Inc. Tax	+Adjust Avg. Lab. Inc. Tax	+Optimal Policy
<i>A. Pct. Ch. in Wealth by Ent. Prod.</i>			
Top 1%	48.3%	52.3%	38.8%
Top 5%	37.0%	45.5%	31.1%
Top 10%	30.4%	41.6%	26.7%
Top 50%	18.4%	36.2%	19.0%
Bottom 50%	3.5%	29.1%	8.0%
<i>B. Pct. Ch. in Wealth by Wealth Group</i>			
Top 1%	37.3%	45.5%	20.6%
Top 5%	25.2%	39.4%	11.4%
Top 10%	21.3%	37.8%	10.4%
Top 50%	15.1%	35.2%	15.3%
Bottom 50%	-12.0%	1.3%	45.7%
<i>C. Pct. Ch. in Wealth by Age</i>			
Age 21-34	-7.4%	-3.1%	1.0%
Age 35-49	0.9%	18.3%	4.5%
Age 50-64	6.1%	32.1%	13.8%
Age 65+	24.0%	40.6%	20.9%

Note: This table reports the effect of the simple fiscal reforms on the distribution of capital. All percentages are computed relative to the benchmark economy. See Table 6 for further details.

quantities increase across the board after the tax reform. In particular, wealth increases by 15.7%, while its quality adjusted counterpart (Q) increases by 31%. The larger increase in Q is driven by the reduction in capital misallocation induced by the shift in taxation from capital income to consumption. This is due to the novel feature of consumption taxation presented in this paper, whereas by taxing consumption, the government can indirectly tax more heavily wealthy unproductive agents at the benefit of productive ones (see more below and in Table 7).¹³

The improvement in efficiency is reflected in an increase in TFP of 4.6%, which follows mechanically from the rise of Q/K . The increase in efficiency also boosts the marginal product of labor (via the real wage) by 10.8. This increase triggers an endogenous effect on average labor income tax rate, that increases by 1.6 percentage points to 15.1%. Furthermore, the wealth effect of increased capital, combined with the income effect of higher real

¹³Taxing consumption rather than capital income is still welfare enhancing in standard OLG models with constant returns to wealth, as it raises capital intensity, thus raising steady-state real wages, leisure and consumption. This channel is stronger than the standard argument of Pigouvian incentive for taxing capital at a positive rate induced by uninsurable labor income risk. However relative to the benchmark economy, the welfare gains in this case are reduced by around 15% in the optimal policy exercise. See Appendix F.

wages, leads to a small decrease in hours worked (-2.3%). This implies that the increase in output (+9.3%) and its private components (i.e., consumption, +10.6% and investment, +21.3%), are mainly driven by higher quality-adjusted capital and not by labor input increases. Furthermore, the reduction in capital misallocation redistributes capital towards high-productivity entrepreneurs, thus pushing some marginal agents out of their backyard-technology businesses (-4.3 p.p.). Those agents still choosing to be entrepreneurs with this partial reform in place are on average more productive than those in the status-quo.

Distributional Effects The effects of this first reform on the allocation of capital are explicitly shown in Table 7. Entrepreneurs with the highest productivity (top 1%) see an increase in their wealth by 48.3%. This is because under the new policy, high-ability entrepreneurs do not pay any tax on the return to their businesses. This creates an incentive to increase their savings, relaxing their financial constraints and further expanding their productive businesses. This effect decreases with lower productivity, with entrepreneurs in the bottom half of the ability distribution enjoying an increase in wealth of only 3.5%.

The same logic holds for households in the top of the wealth distribution, whose composition is skewed towards productive entrepreneurs. In this case, households in the top 1% see an increase in their wealth of 37.3%, while agents in the bottom half of the wealth distribution see a 12% reduction in their wealth. This is because households in the bottom half of the wealth distribution are mainly workers, relying on returns on bonds for consumption smoothing. As the borrowing rate decreases, these agents see a deterioration in the value of their wealth, and given the standard elasticity of savings to interest rate, this implies a lower incentive to save. As Table 6 shows, the reallocation of capital towards the top of the distribution increases the wealth gini by 3.5%, and as a direct consequence, the consumption gini increases by 5.7%. Inequality in labor earnings are almost muted, which is consistent with the small impact of the reform on aggregate and distributional hours.

Finally, Panel C reports the effects by on the distribution of wealth by age. The young (age 21-34) hold less wealth on average (-7.4%). The young are more likely to be wealth poor and most likely a worker (the average age of an entrepreneur is 40). As a result, they suffer from the negative general equilibrium effect on the returns on bonds. The change in wealth monotonically increases along the life cycle—given the increase in efficiency—and peaks in retirement.

Welfare Effects And finally, the overall welfare gain of this reform is large and amounts to 6.9% of unit of consumption per capita per year. Due to higher wealth and consumption inequality, the welfare in “level” increases, however there is a negative welfare effect due to

Table 8 – Welfare Decomposition of Simple Reforms

Variable	Eliminate Cap. Inc. Tax	+Adjust Avg. Lab. Inc. Tax	+Optimal Policy
Cons., CEV_c	5.3%	12.3%	10.2%
Cons. Level, CEV_{cL}	10.6%	14.1%	-2.9%
Cons. Dist., CEV_{cD}	-4.7%	-1.6%	13.5%
Leisure, CEV_h	1.5%	-0.1%	7.1%
Leisure Level, CEV_{hL}	1.2%	-0.3%	7.1%
Leisure Dist., CEV_{hD}	0.2%	0.2%	0.0%
Total Level, $(1 + CEV_{cL})(1 + CEV_{hL})$	12.4%	13.8%	4%
Total Distribution, $(1 + CEV_{cD})(1 + CEV_{hD})$	-4.6%	-1.4%	13%
Total, CEV	6.9%	12.2%	18.0%

Note: This table report decomposes the welfare gains of the simple reforms, using the approach of [Conesa et al. \(2009\)](#). Note that the decomposition is multiplicative – e.g., $1 + CEV = (1 + CEV_c)(1 + CEV_h)$. Further details on partial reforms can be found in [Table 6](#).


redistribution. This finding is mostly driven by the negative welfare effect (-4.7%) of higher dispersion of consumption that the reform brings about, see [Table 8](#). Interestingly, virtually all agents at birth are better off with this policy reform, making it not only welfare improving but also Pareto improving ex-post in both the labor and entrepreneurial ability dimension (see [Figure 1](#)).¹⁴ This is because agents at the top of the distribution who are capital-income rich and likely with high returns, benefit from this reform as they see a decrease in their tax base, and a favourable reallocation of capital. This reallocation of capital increases efficiency, thus boosting real wages and in turn benefiting the agents at the bottom of the distribution who rely more heavily on labor income. Absent of these general equilibrium effects on wages and interest rates, only entrepreneurs at the top of the distribution would gain from this policy, see [Appendix E](#). [Figure 1](#) shows that the welfare gains are evenly distributed along the labor and the entrepreneurial skills, with agents in the top 1 percent of entrepreneurial abilities experiencing only marginally higher gains (7.1% on average).

4.2.2 Partial Reform 2: Eliminate Labor Income Tax on Average

Next, starting from the partial reform 1 (zero tax on capital income and wealth), in the second partial reform, we reduce (optimally) the average labor income tax from 15.1% to its lower bound of zero, but holding fixed the wage tax progressivity at its status-quo value of 0.20. Therefore, the government will subsidize households with low labor income, and tax

¹⁴While we maximize the ex-ante welfare of a newborn, we calculate here the individual welfare at birth, conditional on having a permanent labor productivity type and an initial entrepreneurial ability.

Figure 1 – Welfare Gains from Simple Partial Reforms



Note: This figure reports the welfare gains of the simple reforms by characteristics at birth: labor income type and the initial entrepreneurial productivity.

those with high labor income, but still raise zero revenue in total from the tax. Any lost revenue is recovered by increasing taxes on consumption. In this way, our exercise enables us to study in a clear manner the effects of replacing taxes on labor with those on consumption in isolation from the social insurance benefits of progressive wage taxes. This is different from the analysis of [Nishiyama and Smetters \(2005\)](#), who instead study a scenario where all revenues are raised by taxing consumption but taxes on wages are zero for everyone.

There are various effects that make this reform desirable. First, there is an increase in tax efficiency due to a well known mechanical effect, as consumption has a larger tax base than labor income. Hence, by shifting the burden of taxation from labor income to consumption, a government can meet its budget with lower tax distortions (e.g., [Coleman, 2000](#); [Laczó and Rossi, 2020](#)).¹⁵

¹⁵This can be seen informally in [Table 6](#). The average wedge in the marginal rate of substitution between leisure and consumption ($\tilde{\tau} = 1 - \frac{1-\lambda_t}{1+\tau^c}$) is 0.24 under simple reform 1, and 0.20 under reform 2. Lower

Second, there are benefits due to the life-cycle structure of the economy under consideration (e.g., [Summers, 1981](#)). Since the timing of tax collection is very different under wage and consumption taxes, they have very different implications for saving and consumption profiles. In particular, consumption taxation extracts revenue later in the agent’s lifetime than does wage taxation. This creates an increase in savings earlier in life, which is beneficial as it improves the self-insurance channel created by incomplete markets. As we will show later, this has positive distributional effects. Moreover it is also beneficial for young, high-ability entrepreneurs who use higher savings earlier in life to leverage more capital and expand their business.

Along the same line, another important benefit of the reform stems from the fact that individuals use the real interest rate, r , to compute the present value of taxes paid over their lives, while governments implicitly discount tax revenues at the rate of population growth, n . Therefore, as long as $r > n$ (as it is in the model), the government can raise the same amount of revenues, and reduce the discounted value of the taxes each household pays, by postponing the extraction of taxes. Taxing consumption rather than wage income achieves exactly this goal.

Macroeconomic Effects The beneficial effects of taxing consumption rather than wages are reported in [Table 6](#). Macroeconomic aggregates increase, not only relative to the status-quo, but also relative to the first partial reform. Wealth and its quality-adjusted counterpart raise by 35% and 43%, respectively. The increase over the first simple reform is due, as explained above, by the increase in savings earlier in life. The rise in wealth increases investment (+37.5%), consumption (+14.1%) and output (+14.3%). The increase in savings has a standard detrimental effect on the price of capital (-20.1%) and on the borrowing rate (-1 p.p.). At the same time, higher wealth raise capital intensity, increasing real wages (+13.4%) both relative to the status quo and to the first partial reform. Higher pre and post-tax wages boost hours worked of around 2.5% relative to the first reform, or 0.4% relative to the status-quo.

Distributional Effects As explained above, replacing a tax on wage with one on consumption has far reaching implications for the distribution of wealth (see [Table 7](#)). Relative to the first simple policy, the increase in wealth is more evenly distributed across the population. The reason for this is straightforward. First, since labor income is less concentrated than capital income, eliminating a wage tax directly affects a larger share of the population. Differently, eliminating taxes on capital income has a major effect to a small share of the

numbers indicate lower distortions.

population, those with high entrepreneurial skills. Second, as explained above, this second reform tilts the timing of taxation away from young individuals, who are relatively wealth and income poor, thus benefiting them more.

In this sense, relative to the first reform, eliminating wage tax (keeping fixed progressivity) shifts upward the life-cycle saving profiles of all individuals, not only to those with high entrepreneurial abilities. As a consequence, while wealth increases for all households, in relative terms, agents at the bottom of the distribution increase savings more. The net increase in wealth is around 4 percentage points for entrepreneurs with the highest productivity (i.e., from 48.3% to 52.3%), while wealth increases by more than 25 percentage points for less productive entrepreneurs in the left tail of the distribution. The wealth increase by wealth percentiles tells a similar story. Relative to the first reform, wealthiest individuals in the top 1% increase their savings by 8.2 p.p., while those in the bottom 50 percent by more than 13 p.p. Finally, Panel C presents the effects of the change in the timing of taxation, by age. Saving profiles increase relative to the first reform for all age groups, with the highest relative effect for middle-aged agents (between 50 and 64 years old).

The fact that wealth is less concentrated for individuals at the top of the entrepreneurial ability distribution, decreases, relative to the first reform, both the wealth and consumption gini, with the latter measure decreasing more than 3.2 p.p. (although inequality still increases relative to the status-quo). Similarly, the relatively lower concentration of wealth in the hands of high productivity entrepreneurs is detrimental for TFP, relative to the first reform (-1.6 p.p.), see Table 6. Nevertheless, taken together, the first two simple reforms improve productivity relative to the status-quo by 2%.

Welfare Effects The welfare benefits of replacing wage taxes (on average) with consumption taxes are large and amount to 12.2% relative to the status-quo and 5.3% relative to the first reform. Most of these combined gains come from “level” effects (13.%), while the distributional effects of the first two reforms is still negative (-1.4%). However, for the reasons explained above, relative to the first reform, there is an improvement in the distributional properties of this simple policy (+3.2%). Table 8 presents the decomposition of this welfare benefits in level and distribution effects for consumption and leisure. The main finding is that most of the effects go through changes in consumption levels (+14.1%), while the welfare channel of leisure is almost muted (-0.1%). Finally, as for the first reform, general equilibrium effects on prices and individuals behavior make the welfare gains of this second policy well distributed across labor and entrepreneurial ability levels, see Figure 1. As presented in Appendix E, in a partial equilibrium setting, the joint effects of the first two partial reforms are beneficial only for top ability entrepreneurs. However, relative to the first partial reform,

individuals at the bottom of the distribution experience an increase in welfare. These results imply that the general equilibrium effects are more important for the welfare analysis of the first partial reform.

Taken together, in the first two partial reforms, the planner cannot directly transfer resources between households in order to provide social insurance beyond the status-quo. This means that the aggregate and distributional effects induced by these policies generate exclusively from general equilibrium changes in prices and agents' behavior. In this sense, we extend to consumption taxation the idea of constrained efficient optimal policy in OLG economies with incomplete markets (e.g., [Davila et al., 2012](#); [Krueger et al., 2021](#), [Peterman and Sager, 2022](#)), and show that the welfare gains from taxing consumption are large and uniformly distributed across the population.¹⁶

4.2.3 Partial Reform 3: Move to Optimal Wage Tax Progressivity

In the final step, we abandon constrained efficiency and optimally tilt the progressivity of the wage tax system. Relative to the status quo, the planner can now directly redistribute resources between households. As before, any loss of revenue due to changes in progressivity is compensated by raising consumption taxes. The main original result of this exercise is that within a life-cycle model reproducing income and wealth inequality of the US, adopting consumption taxation calls for high social insurance through redistribution. As we will show clearly in the next section, taxing consumption increases the ability of the planner to redistribute.

Results from this policy experiment can be found in the last column of [Table 6](#). Given that in this last step we replicate optimal policy, one can read the same numbers in [Table 5](#). For this reason we will use the terms “third partial reform” and “optimal policy” interchangeably. The first important result from this exercise is that the planner optimally increases progressivity of wage tax schedule from 0.2 in the status quo to 0.35. This represents a substantial increase. Relative to the second reform, the marginal tax rate for individuals in the top 1% of labor income increases from 48% to 70%.¹⁷

¹⁶Constrained efficiency refers to a policy problem in which the planner cannot directly overcome a friction implied by missing markets (see [Diamond \(1967\)](#)). Here the concept has to be interpreted as relative to the status-quo. In other words, the equilibrium emerging in the second partial reform represents the best outcome the planner can achieve without changing direct redistribution between individuals.

¹⁷We also tried an alternative decomposition of the optimal policy in which we calculate optimal progressivity of labor income with consumption taxation, but keeping the aggregate level of wage tax at the status-quo (see [Appendix F](#)). We find that the optimal progressivity is 0.35, like in the benchmark. This result indicates that the optimal level of social insurance depends exclusively from the characteristics of the model, the planner's welfare function and the presence of consumption taxation. In this alternative scenario, with the average labor income tax is at the status-quo, the marginal tax faced by the richest 1% increases from 38% to 63%.

By providing higher social insurance, the planner compensates for the heterogeneity in labor income due to both transitory and permanent differences in productivity. This is by now a standard result in the literature. First there are no insurance markets (beyond self-insurance) to protect individuals against idiosyncratic transitory fluctuations in wages. With higher progressivity, those who experience good productivity shocks compensate through higher marginal tax rates those individuals who face negative shocks, who in turn receive higher subsidies. Second, there are no markets to insure against the permanent differences at birth (e.g., the veil of ignorance). As for the transitory component, by imposing a progressive tax schedule, the planner can provide insurance on permanent heterogeneity as well. As we will show later in the paper, both transitory and permanent heterogeneity call for higher progressivity in the wage tax schedule.

Macroeconomic Effects Furthermore, relative to the first two reforms, increased progressivity has a distortionary effects on all macroeconomic aggregates. Nevertheless, aggregate wealth and its related variables (i.e., quality-adjusted wealth and investment), still increase relative to the status-quo by 16%, 27% and 18.9%, respectively. Differently, higher progressivity distorts labor supply, leading to a sharp reduction in hours worked relative to the status-quo (-12.1%), without compromising on aggregate output. This is due to the improved allocation of capital. Perhaps more surprisingly, the same distortionary effects of wage taxation reduces private consumption by 2.9%.

Aggregate productivity increases both relative to the status-quo (+3.1%) and, interestingly, also relatively to the second reform (+1 p.p.). This is due to high earners with relatively low entrepreneurial abilities that, given the increase in the marginal wage tax they face relatively to the second partial reform, must reduce the size of their businesses (more on this below). Moreover, the higher supply of capital and its redistribution towards the lower end of the population, put downward pressure on the price of capital, which decreases relative to the status quo (-20.9%), as well as relative to either of the first two reforms. Importantly, the reduction in hours worked, together with the increase in capital intensity and its productivity, pushes real wages upward, with an increase relative to the status quo (+14.1%) and to the first two reforms (0.7%).

Distributional Effects The substantial increase in the wage tax progressivity reduces wealth inequality both relative to the status-quo (-2.2%) and also with respect to the first two reforms reform (-4.1%). The pre-tax earnings gini is not affected by the policy, signaling that the reduction in hours worked is spread evenly across population (i.e., the distortionary effects of higher progressivity reduce labor supply evenly over the whole distribution of

abilities). Perhaps most importantly, the consumption gini decreases by 10.7%, implying that the cross-sectional dispersion of private consumption falls significantly.

These results are reflected in the the distribution of capital (see Table 7). First of all, Panel A shows that relative to the status quo, in the optimal policy, equilibrium capital accumulates more at the top of the entrepreneurial ability distribution. However, this effect is reduced relative to the the first two partial reforms. This is mainly driven by those individuals with both high labor ability and entrepreneurial productivity who now face an increase in their wage marginal tax rate and, as such, they have less resources to invest into their businesses. Panel B shows that the substantial increase in wage progressivity boost the wealth in the bottom half of the distribution, allowing them to smooth consumption more efficiently thus reducing consumption gini as presented in the above paragraph. Interestingly, the increase in wealth is U-shaped, decreasing up to the top 5% of the distribution and then increases again for the top 1%. This is mainly driven by individuals with high entrepreneurial skills who are over-represented in the top 1% of the wealth distribution. Finally, Panel C shown that the combined effect of the three reforms increases savings across the whole life-cycle, with the highest effect for retirees. Nevertheless, relative to the second reform, young agents (i.e., those who are more likely to be workers with low income) gain the most in terms of wealth.

Welfare Effects Table 8 reports a few interesting results regarding the decomposition of welfare under the optimal policy. First of all, relative to the status quo, the level effect of consumption is negative (-2.9%), mirroring the overall decrease in this component. This result is even more evident if compared with the large increase in consumption level that the first two partial reforms brought about (+14.1%). These numbers shed light on the optimal redistribution calls for by our model. The planner is willing to sacrifice a large size of private consumption in order to reduce its unequal distribution. Relative to the status quo, the overall welfare benefits from redistribution are very large (i.e., 13.5%). However this is only part of the story. The planner can maintain the overall size of the economy, as measured by aggregate output, with a much lower labor input (see Table 7). This, in turn, increases leisure in equilibrium, thus boosting utility. The contribution of the leisure channel in the overall welfare gains is substantial (+7.1%) and uniformly distributed over the population (i.e., the gains from the distribution component in leisure is zero). The combined effects of higher redistribution in consumption and higher level in leisure translates into gains in level (+4%) and redistribution (+13%). Not surprisingly, the large increase in progressivity substantially increases welfare for individuals in the bottom of the distribution and decreases those of the agents in the top 10% of the income distribution (see Figure 1). Perhaps more

surprisingly, for each labor ability category, those with the highest entrepreneurial ability enjoy the largest benefits of increased progressivity. Overall, around 88% of newborns are better off in the new stationary equilibrium under the optimal policy.

Robustness In Appendix F we describe a number of robustness checks. These include: i) a model without return heterogeneity; ii) a model without transitory earnings risk; iii) a model with decreasing returns to scale; iv) a model with rich intergenerational links and joint distribution between labor and entrepreneurial abilities; v) an alternative second partial reform in which we fix aggregate level of average wage tax and we tilt progressivity optimally.¹⁸

4.3 Optimal Policy with Fixed Consumption Tax

In order to isolate precisely the contribution of taxing consumption on optimal fiscal policy and its equilibrium variables, we conduct the same exercise as in the benchmark, but we fix the consumption tax rate at its status-quo value of 7.5%. In this way we can isolate precisely the relative merits consumption taxation vis-a-vis other tax components. We decompose the optimal policy when consumption tax cannot be adjusted optimally into steps that are similar to those in the benchmark exercise (see Table 6 for further details). Results from these experiments can be found in Tables 9-11. The main finding from this experiment is that when the consumption taxation cannot be used beyond its status-quo value, then taxing wealth is optimal. Specifically, the planner decides to set capital income tax to zero, tax wealth at around 3%, raise a substantial amount of revenues from earnings and increase labor income tax progressivity. The welfare gains of this policy are large and amount to around 9.8% of CEV. This policy is beneficial both for productivity (i.e., TFP increases by 6.6% relative to the status-quo), and inequality, as measured by the consumption Gini. These findings mirror those presented in Guvenen et al. (2022b).¹⁹ Results from this experiment are presented as follows.

Partial Reform 1 with Fixed Consumption Tax We start from the partial reform in which we replace in a revenue-neutral manner capital income taxation with wealth tax. First, qualitatively, this partial reform produces similar mechanism as the corresponding one

¹⁸We also experiment with other various tax functions, such as a non-linear capital income taxation, but the results were so similar to the benchmark, we decided to not report it for sake of brevity.

¹⁹There are two main differences between the current study and Guvenen et al. (2022b). First, we consider progressive wage taxation rather than linear taxes. Second, in our model, only a fraction of agents are entrepreneurs, and this fraction is one of our targets in the estimation exercise. Differently, in Guvenen et al. (2022b) all households run businesses.

in Table 6, when we replaced capital income tax with consumption tax. As in that case, the policy is expansionary on macroeconomic aggregates, improves TFP and real wages, and boosts inequality measures like the wealth and consumption gini. At the same time, it is welfare improving, with gains of around 6.3%, with a positive effect on level and a negative on distribution. Comparing stationary equilibria, all households benefit from this reform, relative to the status-quo. The reason for these findings is by now well understood: under capital income taxation and heterogeneous entrepreneurs, high returns businesses pay relatively more taxes than low productivity ones. Differently, under wealth taxation, agents who have similar assets pay the same taxes irrespective of their productivity levels. This in turn shifts the relative burden of taxation away from high productivity entrepreneurs, thus increasing their incentive to save, relaxing their financial constraints and enlarging their productive businesses.

However, relative to consumption taxation, a wealth tax erodes the principle, which is detrimental on the aggregate level of wealth, particularly for agents with low entrepreneurial abilities, who are most likely workers. As a result, the reallocation of capital towards entrepreneurs with the highest productivity is relatively larger under wealth tax than under consumption taxation (see Table 10). This has a stronger effect on TFP (as now top entrepreneurs have relatively more capital) but a negative effect on the saving rate of agents in the bottom half of the wealth distribution. In turn, the lower accumulation of aggregate savings boosts less (relative to consumption taxation) capital intensity, thus the increase in real wages is relatively more modest (6.6% vs 10.8%).

In summary, shifting the burden of taxation from capital income to consumption leads to marginally higher welfare gains than shifting taxation from capital income to wealth (6.9% vs 6.3% of CEV). First, taxing wealth deteriorates the principal and hence the increase in aggregate savings is lower than under consumption taxation. This leads to a lower increase in real wages, thus benefiting less those households at the bottom of the distribution, mostly relying on earnings. Second, as wealth is more concentrated than consumption (and disproportionately so among entrepreneurs), the burden of wealth tax falls relatively more on high-productivity entrepreneurs, thus leading to lower welfare benefit for those at the top of the entrepreneurial distribution (see Figure 2).

Partial Reform 2 with Fixed Consumption Tax Then starting from the first partial reform (zero tax on capital income and consumption taxation at its status-quo), in the second partial reform, we analyze the optimal composition of wealth vs. wage taxation under the constraint that the planner does not change the progressivity of the system. In other words, like in the second partial reform with consumption taxation, the government cannot modify

Table 9 – Simple Reforms with Consumption Tax Kept at the Status-Quo

Variable	Status-Quo	Eliminate Cap. Inc. Tax	+Adjust Avg. Lab. Inc. Tax	+ Optimal Policy
<i>Tax Instruments:</i>				
Consumption Tax	7.5%	7.5%	7.5%	7.5%
Avg. Labor Income Tax	13.5%	14.6%	9.1%	13.2%
Labor Tax Progressivity	0.20	0.20	0.20	0.30
Capital Income Tax	25%	0%	0%	0%
Wealth Tax	0%	1.5%	3.2%	3.0%
<i>Aggregate Quantities:</i>				
Wealth		1.5	-7.3	-17.4
Quality-Adj. Wealth		18.7	10.0	-1.3
Hours		-1.3	0.2	-9.1
Output		6.0	4.0	-6.1
Consumption		8.8	7.9	-5.1
Investment		6.9	-2.6	-13.2
<i>Productivity:</i>				
TFP		5.8	6.4	6.6
Entrepreneurial Rate (Δ p.p.)		-5.1	-5.4	-5.5
<i>Prices:</i>				
Price of Capital		-10.7	-5.4	-4.9
Wages		6.6	3.2	2.8
Borrowing Rate (Δ p.p.)		0.5	1.6	1.9
<i>Inequality:</i>				
Wealth Gini		5.6	7.6	3.6
Earnings Gini		0.5	0.4	0.3
Consumption Gini		4.0	1.6	-7.7
<i>Welfare:</i>				
CE Welfare Gain		6.3	7.1	9.8
<i>Level</i>		8.5	7.1	0
<i>Redistribution</i>		-3.0	0.0	9.8
Pct. of Newborns Better Off		100	99.9	87

Note: This table reports a sequence of revenue-neutral simple tax reforms in which the wealth tax is used to raise any additional revenue. In the first reform, the capital income tax is replaced with a wealth tax. As it is infeasible to reduce the average labor income tax, the second reform further reduces the average labor income tax rate and increases the wealth tax to the optimal levels, constraining the labor income tax to stay fixed at its benchmark value. In the final reform, the economy moves to the optimal progressive labor income and wealth tax. In all cases, we hold fixed the consumption tax rate at its benchmark value. All variables (except for the tax instruments) are reported as the percentage or percentage point change relative to the benchmark economy.

directly the transfers between agents. This means that we keep the wage tax progressivity constant at 0.2.

There are various reasons why the planner wants to shift at least part of the tax burden from wages to wealth. First of all, as in standard OLG settings, by doing so, the government can twist the timing of tax collection. As young agents are relatively wage rich and wealth poor, taxing the latter implies postponing tax extraction. As for the case of consumption taxation, postponing tax extraction might be beneficial as long as the real interest rate is higher than the population growth rate. However, in stark difference from consumption taxation, the planner discourages savings by taxing wealth. This has a detrimental effect in a context with incomplete markets, as it decreases self-insurance and therefore it induces less efficient consumption smoothing.

Second, by taxing wealth rather than earnings the planner can shift the tax burden towards a highly concentrated tax base (as wealth is more unequal than earnings), thus indirectly redistributing from rich to poor households. However, this distributional goal has to be balanced against the damage that a higher wealth tax brings about on the principal (i.e., on aggregate savings). In the constrained efficient allocation, the planner taxes wealth at 3.2% and imposes an average wage tax of 9.1%. This is different than the constrained efficient policy analyzed in the benchmark exercise, where the planner decides to raise all revenues from taxing consumption. As we will describe below, this is the manifestation of a much stronger distortionary nature of wealth taxation, which in turn compromises its fiscal space and welfare benefits.

The overall effects of higher wealth tax sees an overall contraction of aggregate savings (-7.3%), although an increase in its quality-adjusted measure (+10%), see Table 10. The combined effects of these partial reforms are beneficial for aggregate output (+4%) and private consumption (+7.9%), but creates a small decrease in private investment (-2.6%), which directly follows from the decrease in aggregate savings. The overall effect on TFP is positive both relative to the status quo (+6.4%) and relative to the first partial reform (+0.6%). This is because a higher wealth tax pushes some marginal entrepreneurs out of business, thus reallocating capital towards higher-productivity businesses, mechanically increasing TFP. Table 10 shows this effect clearly. Under the second partial reform, only agents in the top 10% of the entrepreneurial ability distribution sees a gain in their wealth share- relative to the status-quo, while agents in the top 50% of the distribution (and below) experience a decrease in their savings (always relative to the status-quo). The same effects are mirrored by wealth shares in Panel B.

Turning to the effects of this reform on equilibrium prices, we find that the higher quality-adjusted wealth increases real wages relative to the status-quo (+3.2%), however, it decreases

real wages relative to the first partial reform. This is because the substantial increase in the wealth tax erodes the principle, thus reducing capital intensity. Along the same line, the lower level of capital, coupled with higher quality-adjusted capital, implies that the price of capital decreases (-5.4%), while the pre-tax borrowing rate jumps by 1.6 p.p.

By increasing the reallocation of capital towards highly-productive entrepreneurs, the second partial reform further boosts the wealth gini, and by increasing the tax base of wealth taxation, further shifts the burden of taxation to the top of the distribution. As a result, the increase in consumption gini is lower relative to the first partial reform. The combined welfare benefits of the first two partial reforms is large and amount to 7.1% CEV. All benefits come from level effects, with virtually all agents gaining from this reform. Interestingly, the higher concentration of the wealth tax base makes agents with the highest entrepreneurial abilities gain the least from these two partial reforms (see Figure 2). While the welfare gains of the constrained efficient policies are substantial, they are smaller relative to those in the benchmark exercise, where all revenues were raised with consumption taxation. Relative to that scenario, wealth tax is more distortionary, mainly on wealth. This has negative welfare effects along the ability distribution and along agents' life-cycles. Overall, the benefit of raising all revenues from consumption (without modifying redistribution relative to the status-quo) increase the welfare gains by slightly more than 40%, or 5.1 p.p., in CEV.

Optimal Policy with Fixed Consumption Tax Finally, starting from the second reform, we optimally tilt the progressivity of wage tax. The optimal progressivity of the labor tax schedule is higher than in the status-quo, with an increase in relative progressivity of around 30% (i.e., $\frac{PW_{\text{Optimal Policy}(\tau_c = 0.075)}}{PW_{\text{Status-Quo}}} = 1.30$), whereas the marginal tax rates faced by the richest 1% are around 70% and the poorest 50% around 30.4%. Therefore the planner desires to increase redistribution and social insurance and balances this desire against the strong distortionary effects of higher progressivity – i.e., labor supply (-9.1%), capital accumulation (-17.4%) and the overall size of the economy (-6.1%). While the size of the economy shrinks, it operates more efficiently with higher TFP (+6.6%) and higher real wages (+2.8%).

Not surprisingly, the increase in progressivity greatly reduces inequality in the post-tax earnings gini (-4.5%) and in the consumption gini (-7.7%) and more than halves the increase in wealth gini relative to the first two partial reforms (3.6% vs. 7.6%). The marginal gains of optimal progressivity are substantial and amount to a 2.7 p.p. increase relative to the first two partial reforms. This pushes the overall welfare gains to 9.8% CEV. All benefits from the optimal policy come from distributional gains, mainly from consumption (see Table 11), and the overall level effect of the optimal policy is close to zero. This is because the increase in welfare deriving from higher leisure (+5.4%) compensates for the negative effects of lower

aggregate consumption (-5.1%), see Table 11. Not surprisingly, individuals at the bottom of the distribution are those who gain the most from the optimal policy, while most individuals (87%) are better off with the optimal policy in place (see Figure 2).

Relative to the benchmark policy experiment in which the planner taxes consumption optimally, the welfare gains are reduced by around 50%. This large loss of welfare derives both from distributional benefits (13.5% vs 7.1%) and levels effects (4.0% vs -0.1%). Also, taxing consumption optimally is beneficial along the whole distribution of labor and entrepreneurial abilities (see Figure 2).

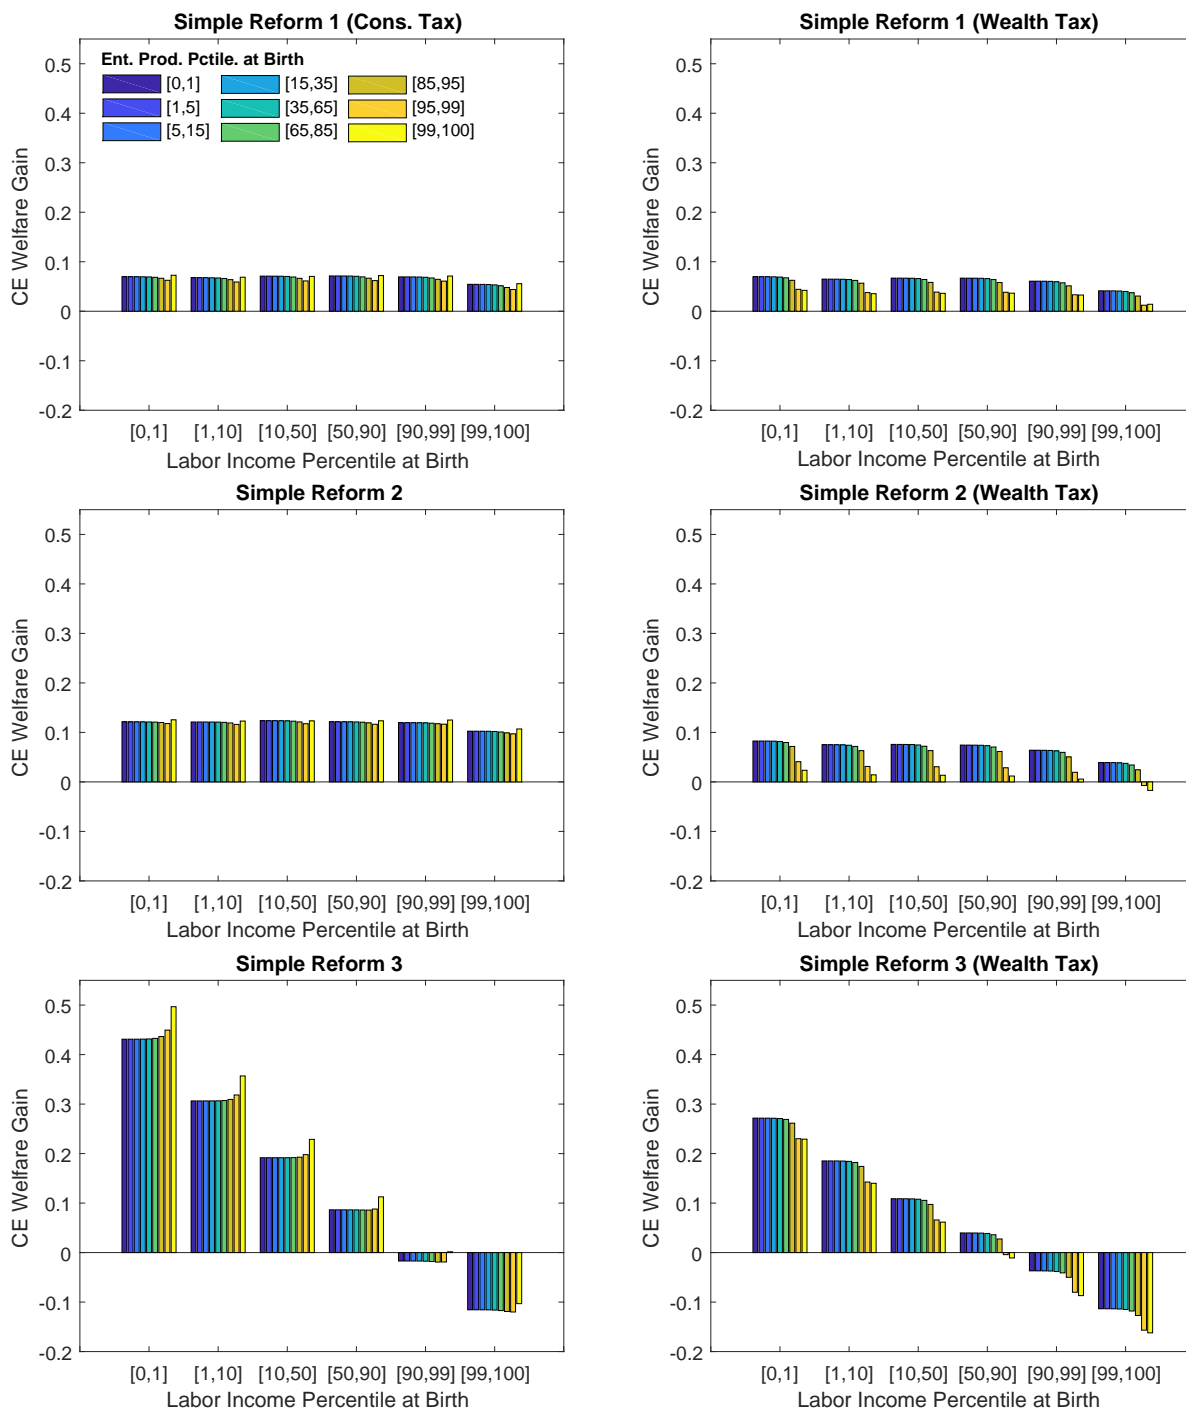
First, taxing consumption rather than capital allows the policymaker to reallocate capital efficiently and at the same time, change the tax base from a highly concentrated tax base (capital income) to a less concentrated one (consumption). Both these effects improve efficiency and reduce distortions, thus boosting real wages more than when τ_c is constrained at the status quo. In turn, this benefits individuals along the whole income and wealth distribution. Second, the planner has the fiscal space to fully substitute wage taxes with consumption taxes. This allows the planner to postpone the timing of tax collection, allowing young agents to increase their savings, thus providing better self-insurance in an incomplete markets economy. At the same time, high-productivity workers who also happen to have large entrepreneurial abilities can further expand their productive businesses. This large fiscal space is lost when τ_c is kept and its status-quo and wealth tax is used instead, leading to much lower welfare gains.

Finally, the efficiency of optimal consumption taxation allows the policymaker to provide stronger redistribution – i.e., the optimal progressivity under consumption taxation is around 23% higher than in the alternative scenario analyzed here ($\frac{PW_{\text{Optimal Policy}}}{PW_{\text{Optimal Policy}(\tau_c = 0.075)}} = 1.23$). Interestingly, the difference between progressivity is implemented via lower wedges at the bottom of the distribution (17.4% in the benchmark and 30% here), rather than with higher taxes at the top (in both cases around 70%). The resulting optimal redistributive policy leads to higher welfare gains for agents at the bottom of the distribution and lower welfare losses for those with very top labor and entrepreneurial abilities.

4.4 Transitional Dynamics

In this section, we focus once again exclusively on the benchmark case (i.e., where consumption can be freely taxed), and consider the effects of our fiscal reforms by studying short-run transitional dynamics. This is important as the effects of a policy on newborns in a stationary equilibrium (i.e., in the long-run) might be radically different from the effects on individuals alive when the policy is initially implemented. This is particularly true in the

Figure 2 – Welfare Gains from Simple Reforms with Wealth Tax



Note: This figure reports the welfare gains from the simple reforms compare when a consumption tax is used to raise lost revenue (left panels) and when a wealth tax is used instead (right panels).

Table 10 – Effect of Simple Reforms on Allocation of Capital, Consumption Tax at the Status-Quo

Variable	Eliminate	+Adjust	+Optimal Policy
	Cap. Inc. Tax	Avg. Lab. Inc. Tax	
<i>A. Pct. Ch. in Wealth by Ent. Prod.</i>			
Top 1%	38.5%	30.4%	18.3%
Top 5%	22.8%	12.0%	0.6%
Top 10%	16.0%	5.4%	-5.5%
Top 50%	3.9%	-5.4%	-15.4%
Bottom 50%	-11.5%	-20.0%	-28.5%
<i>B. Pct. Ch. in Wealth by Wealth Group</i>			
Top 1%	25.6%	17.0%	-0.1%
Top 5%	13.3%	6.0%	-12.0%
Top 10%	9.2%	2.0%	-15.0%
Top 50%	1.0%	-7.6%	-18.3%
Bottom 50%	-45.3%	-71.4%	-45.0%
<i>C. Pct. Ch. in Wealth by Age (at Birth)</i>			
Age 21-34	-42.0%	-66.2%	-59%
Age 35-49	-19.2%	-36.0%	-36.3%
Age 50-64	-7.2%	-13.7%	-23.7%
Age 65+	10.1%	1.6%	-10.5%

Note: This table reports the effect of the simple fiscal reforms on the distribution of capital. All percentages are computed relative to the benchmark economy. Consumption tax rate kept at its status-quo value of 7.5%. See Table 6 for further details.

Table 11 – Welfare Decomposition of Simple Reforms, Consumption Tax at the Status-Quo

Variable	Eliminate	+Adjust	+Optimal Policy
	Cap. Inc. Tax	Avg. Lab. Inc. Tax	
Cons., CEV_c	5.3%	7.1%	4.2%
Cons. Level, CEV_{cL}	8.8%	7.9%	-5.2%
Cons. Dist., CEV_{cD}	-3.2%	-0.7%	9.9%
Leisure, CEV_h	0.9%	0.0%	5.4%
Leisure Level, CEV_{hL}	0.7%	-0.1%	5.4%
Leisure Dist., CEV_{hD}	0.2%	0.1%	0.0%
Total Level, $(1 + CEV_{cL})(1 + CEV_{hL})$	9.6%	7.7%	-0.1%
Total Distribution, $(1 + CEV_{cD})(1 + CEV_{hD})$	-3.0%	-0.6%	9.9%
Total, CEV	6.3%	7.1%	9.8%

Note: This table report decomposes the welfare gains of the simple reforms, using the approach of Conesa et al. (2009). Note that the decomposition is multiplicative – e.g., $1 + CEV = (1 + CEV_c)(1 + CEV_h)$. Further details on partial reforms can be found in Table 6.

context of our analysis about shifting the tax burden towards consumption. For example, since consumption takes place later in life than labor income, a switch to consumption taxation would transfer the tax burden towards the elderly. Thus, older agents (e.g., retirees) alive at the time of the reform pay relatively more taxes, while younger households and subsequent newborns pay less by having their tax payments deferred to older ages. This provides a substantial increase in the welfare of young and future individuals, and potentially a utility loss for some generations alive at the time of the switch. This type of consideration about the intergenerational redistribution of fiscal policy gives rise to a number of implementation concerns.

While solving for the optimal path of taxes with transition is beyond the scope of this paper, here we study an equilibrium that arises during the transition to the new long-run equilibrium. Specifically, we implement this exercise by fixing consumption tax to its steady-state values. Moreover, we allow the government to run public debt along the transition, and adjust the level of the income tax (via λ_t) only during the transition, so that the government budget constraint is balanced in the new stationary equilibrium. It is important to note that the long-run steady state is isomorphic to that presented in Table 6. Thus, while we allow for debt along the transition path, its value continues to be zero in the long run. Results from this exercise are presented in Table 12.²⁰

We begin with the first partial reform in which we substitute capital income tax with consumption tax. Newborns and young workers (i.e., less than 40 years of age) gain the most from this reform as it allows them to accumulate more capital earlier in life and, in relative terms, postpones the extraction of their tax burden. These gains decrease for middle-aged agents. While they enjoy lower taxes on capital both as workers and also as entrepreneurs, they did not accumulate enough capital to optimally face the higher price of consumption during their retirement years, which they will experience in their near future. Not surprisingly, the smallest benefits of the first reform are experienced by the elderly. While they enjoy higher net income from their savings, these are decreasing during the final years of their lives. Most importantly, in relative terms, consumption taxes are particularly harmful for the elderly alive at the time of the reform. While in aggregate, agents aged 65+ have a welfare gains of 0.7% in CEV, some agents in the final years of their lives experience a small welfare loss. On impact, the aggregate welfare gains of this partial reform is around 2.7% of CEV, with 92.3% of the agents benefiting on impact.

With the second partial reform, the intergenerational effects are more apparent. This is

²⁰We also experimented with the alternative case in which the planner balances its budget over the transition path by adjusting consumption taxes rather than wage taxes. Results are, for all practical purposes, unchanged.

Table 12 – Short-Run Welfare Effects of Fiscal Reforms

Variable	Eliminate Cap. Inc. Tax (Reform 1)	+Adjust Avg. Lab. Inc. Tax (Reform 2)	+Optimal Policy (Reform 3)	+Optimal Policy (Pareto)
Individuals on Impact				
Newborn	3.4	7.9	11.0	6.9
Age 22-30	3.6	7.7	8.4	6.0
Age 31-40	3.6	6.3	4.2	3.6
Age 41-50	3.4	3.6	-0.6	1.2
Age 51-64	3.0	-1.4	-8.5	0.6
Age 65+	0.8	-8.6	-17.5	0.4
<i>Average</i>	2.7	0.9	-3.6	2.3
<i>Pct. Support</i>	92.3	62.0	40.1	100
Future Newborns				
+10 years	5.5	10.5	12.4	6.7
+20 years	6.2	11.6	13.1	6.8
+30 years	6.4	11.8	13.4	6.9
<i>Long Run</i>	6.9	12.2	18.0	18.0

Note: This table reports the welfare gains on impact and along the transition from implementing the simple reforms.

because, relative to eliminating only capital income taxation, the wage tax base is larger and entirely concentrated away from retirees, who instead pay exclusively taxes on consumption. For this reason, it is not surprising that younger agents gain a lot from this reform, as it postpones tax extraction and allows them to increase savings and smooth consumption more efficiently. At the same time, individuals older than 54 years of age suffer from this reform, both relative to the status-quo and to the first partial reform. This is because their remaining life-time tax liabilities increase substantially. As such, by raising consumption taxation from 12.2% to almost 26%, the planner imposes a large shift in the tax burden from the young to the elderly. On the balance, the overall welfare gains for the generations alive at the time of the reform is positive and amount to around 1% in CEV terms. A large majority of agents (62.0%) would favor this reform.

The transition analysis of the second partial reform in which all revenues are raised by taxing consumption, leads to different conclusions to [Nishiyama and Smetters \(2005\)](#). In particular, they find that shifting all revenues from income to consumption is detrimental not only for the welfare of the current generation, but it also creates an efficiency loss across generations. There are a number of differences between the settings we study. First of all, in their paper return to savings is homogeneous across individuals, while in our economy

returns are heterogeneous and stochastic. As explained in details throughout the paper, this feature creates a novel channel that increases the advantages of consumption taxation. Second of all, while in our setting we collect all revenues via consumption taxation, we maintain social insurance for idiosyncratic risk by fixing the progressivity of wage tax at the status-quo, while in [Nishiyama and Smetters \(2005\)](#), social insurance is eliminated as the wage tax is zero for all individuals. In this sense, a common message of both studies is that consumption tax needs to be implemented without compromising on social insurance, which is exactly what partial reform 2 does.²¹

Then, we analyze the transitional dynamics of implementing the optimal policy. Relative to the second partial reform, the planner increases the progressivity of the wage tax schedule to its optimal level (i.e., $\tau_l = 0.35$). The intergenerational consequences are even more dramatic for this case. While the increase in progressivity increases the ex-ante welfare gains of newborns (+11% in CEV), who benefit from higher social insurance and a delayed extraction in their tax burden, anyone older than 41 years of age on impact (i.e., middle-aged individuals) experience a loss. This loss becomes very large for retirees (-17.5% in CEV). This is due to the combined effect of a higher consumption tax rate that this policy implies (from 25.8% to 30.4%) and higher progressivity, which is harmful for those middle-aged workers who are at the peak of their earning potentials. In aggregate, those alive at the time of the reform experience a welfare loss of around 3.5% in CEV. Given these findings, a natural question is whether our proposed policy remains beneficial once transitional dynamics are taken into account.

Evaluating precisely the intergenerational effects of our proposed reform is not obvious, as there is no uncontroversial welfare criterion that aggregates utility of different generations. Therefore, we use the idea of the Lump-Sum Redistribution Authority (LSRA) described in [Auerbach et al. \(1983\)](#) and subsequent contributions. This allows us to consider transitional dynamics and, at the same time, deviate as little as possible from our benchmark exercise. In a nutshell, this authority uses lump-sum taxes and transfers to keep cohorts born before our tax reform is announced at least at their status-quo level of welfare. Then it asks whether the welfare of all cohorts born after the announcement date are better (or worse) off than in the status-quo, once the cost of compensating those agents who initially lost from the policy are fully repaid via higher distortionary taxes. To this end, the solution of our benchmark model is modified to solve for the economy's general equilibrium transition path consistent with the behavior of the standard government budget constraint as well as the fiscal consequences deriving from the activity of the LSRA. The working of this is the following.²²

²¹See also [Leung and Poschke \(2022\)](#), who analyze the effects of increasing consumption taxation in isolation, within a model with bequests and efficient distribution of capital.

²²Appendix [G](#) presents an alternative exercise in which we impose that any optimal long-run policy has

Assume that our optimal policy is announced at the beginning of period 1. The LSRA makes a lump-sum transfer (tax if negative) to each living household to bring its expected remaining lifetime utility back to (at least) its pre-reform level in the status-quo economy. However, we do not allow the LSRA to put agents in debt. This set of tax/transfers imposes a fiscal need to the government, who finances it via issuing public debt, thus crowding out part of the private capital accumulation that occurs in the benchmark economy without LSRA. Moreover, given that public debt is zero in the long-run, future newborns will have to pay higher distortionary wage taxes (with a consumption tax, results are virtually identical) in order to repay the outstanding debt and its relative interest along the transition. Obviously, this policy implies higher taxes for future newborns. If, given the increase in the tax burden and all other general equilibrium spillovers on prices and individuals' behaviors, the welfare of newborns along the transition remains higher (lower) under the new policy relative to the status-quo, we conclude that our policy is beneficial (detrimental) from an interim point of view. In practical terms, we are checking whether our proposed policy produces enough resources to increase welfare to all future newborns, once the potential losers are fully compensated and the cost of this compensation is entirely taken into account. Therefore, we are analyzing whether is feasible for our proposed policy to improve welfare in a Pareto sense.²³

Results from the last column of Table 12 shows that after compensating the losers, our policy reform increases welfare for all newborns.²⁴ Clearly, the welfare gains of newborns is reduced relative to the same reform but without LSRA, as now their tax burden includes also the transfer to compensate those cohorts born before the policy announcements. The welfare gains of newborns 20 and 50 years after the reform is 6.8% and 10%, respectively. These findings imply that our tax proposal, where all revenues are raised via consumption taxation and social insurance is provided by a progressive wage tax, produces enough resources to increase welfare also along the transition path.

5 Conclusion

This paper revisits the age-old policy quest of efficient redistribution in unequal societies. Our main point is that adopting simple linear consumption taxes enables the government

to be voted in by a at least 50% of the current cohorts. This policy retains around 90% of the long-run gains of the benchmark exercise.

²³Note that this formulation is slightly different to the original [Auerbach et al. \(1983\)](#). In their paper, the LSRA cares only about efficiency, in a compensating variation fashion. In our case, in order to be consistent with our benchmark optimal policy exercise in Section 4.1, we also have a strong redistribution motive due to the maximization of ex-ante welfare of newborns and we measure equivalent variation.

²⁴Note that most age groups experience now a welfare gains. This is because we do not allow the LSRA to put agents in debt, which would be required if wanted to have a zero net impact on current cohorts.

to achieve a better equality-efficiency trade-off. We show this result within an estimated life-cycle model that replicates the high concentration of income and wealth in the United States. We show that the optimal policy calls for a consumption tax of around 30% and a labor income tax that raises no revenue on average. Nevertheless, the optimal policy calls for a substantial increase in labor income tax progressivity, so that the poor would receive higher subsidies and the rich would face higher marginal tax rates. Meanwhile, as long as the government is able to tax consumption optimally, it is not optimal for the government to tax wealth or capital income.

The welfare gains from this reform are large (18%) and while more than two-thirds of these gains are due to redistribution, a big chunk of these gains are reflected in improvements in productivity. By using consumption taxes to raise revenue efficiently, the government is better able to use the labor income tax solely for redistribution. Moreover, we show how this optimal policy could be implemented in a sequence of simple reforms. The first two reforms would achieve large gains (12%), virtually all individuals would benefit from it and improving without any distributional consequences (i.e., the benefits are roughly equally distributed). It is only in the last reform where the progressivity of the labor income tax is increased where high ability households are made worse off, but increase the ex ante welfare gains to 18%. One interesting interpretation of the benefit of taxing consumption, is that even before entering into the debate on the appropriate level of redistribution in the economy, our paper demonstrates the substantial gains that can be made with consumption taxes favoring in the long-run all individuals.

Finally we show that benefits from taxing consumption are robust to transitional concerns. In particular we show that from an intergenerational point of view, our reform generates enough resources so that all newborns along the transition gain even after compensating all those cohorts born before the reform who potentially lose from it.

There are several avenues for future research. First, one could extend the model to include human capital in both labor and entrepreneurial activity (e.g., see [Badel et al., 2020](#)). While this is expected to increase the distortionary effects of progressive wage taxation, it could further increase the role of consumption taxation as a mean to increase allocative efficiency. Second, one can extend the model to include tax avoidance (see [Di Nola et al., 2021](#)). Given consumption is more difficult to hide than assets, this might create a further argument in favor of taxing consumption.

References

- Altig, David, Alan J Auerbach, Laurence J Koltikoff, Kent A Smetters, and Jan Walliser. 2001. “Simulating Fundamental Tax Reform in the United States.” *American Economic Review* 91 (3):574–595.
- Altig, David, Alan J Auerbach, Laurence J Kotlikoff, Elias Ilin, and Victor Ye. 2020. “Marginal Net Taxation of Americans’ Labor Supply.” Tech. rep., National Bureau of Economic Research.
- Altonji, Joseph G. and Lewis M. Segal. 1996. “Small-Sample Bias in GMM Estimation of Covariance Structures.” *Journal of Business & Economic Statistics* 14 (3):353–366.
- Auerbach, Alan J and Laurence J Kotlikoff. 1987. “Evaluating Fiscal Policy with a Dynamic Simulation Model.” *The American Economic Review* 77 (2):49–55.
- Auerbach, Alan J., Laurence J. Kotlikoff, and Darryl Koehler. 2022. “U.S. Inequality and Fiscal Progressivity: An Intragenerational Accounting.” *Journal of Political Economy* 0:null.
- Auerbach, Alan J., Laurence J. Kotlikoff, and Jonathan Skinner. 1983. “The Efficiency Gains from Dynamic Tax Reform.” *International Economic Review* 24 (1):81–100.
- Bach, Laurent, Laurent E. Calvet, and Paolo Sodini. 2020. “Rich Pickings? Risk, Return, and Skill in Household Wealth.” *American Economic Review* 110 (9):2703–2747.
- Badel, Alejandro, Mark Huggett, and Wenlan Luo. 2020. “Taxing Top Earners: A Human Capital Perspective.” *The Economic Journal* 130 (629):1200–1225.
- Benhabib, Jess and Alberto Bisin. 2018. “Skewed Wealth Distributions: Theory and Empirics.” *Journal of Economic Literature* 56 (4):1261–91.
- Benhabib, Jess, Alberto Bisin, and Mi Luo. 2019. “Wealth Distribution and Social Mobility in the US: A Quantitative Approach.” *American Economic Review* 109 (5):1623–1647.
- Benhabib, Jess, Alberto Bisin, and Shenghao Zhu. 2011. “The Distribution of Wealth and Fiscal Policy in Economies with Finitely Lived Agents.” *Econometrica* 79 (1):123–157.
- Bhandari, Anmol, Serdar Birinci, Ellen R. McGrattan, and Kurt See. 2020. “What Do Survey Data Tell Us about US Businesses?” *American Economic Review: Insights* 2 (4):443–458.
- Boar, Corina and Virgiliu Midrigan. 2019. “Markups and Inequality.” Working Paper.

- . 2022. “Should We Tax Capital Income or Wealth?” *The American Economic Review: Insights* .
- Cagetti, Marco and Mariacristina De Nardi. 2006. “Entrepreneurship, Frictions, and Wealth.” *Journal of Political Economy* 114 (5):835–870.
- Coleman, Wilbur John II. 2000. “Welfare and Optimum Dynamic Taxation of Consumption and Income.” *Journal of Public Economics* 76 (1):1–39.
- Conesa, Juan Carlos, Sagiri Kitao, and Dirk Krueger. 2009. “Taxing Capital? Not a Bad Idea after All!” *American Economic Review* 99 (1):25–48.
- Correia, Isabel. 2010. “Consumption Taxes and Redistribution.” *American Economic Review* 100 (4):1673–1694.
- Correia, Isabel, Emmanuel Farhi, Juan Pablo Nicolini, and Pedro Teles. 2013. “Unconventional Fiscal Policy at the Zero Bound.” *American Economic Review* 103 (4):1172–1211.
- Davila, Julio, Jay H Hong, Per Krusell, and José-Víctor Ríos-Rull. 2012. “Constrained Efficiency in the Neoclassical Growth Model with Uninsurable Idiosyncratic Shocks.” *Econometrica* 80 (6):2431–2467.
- Di Nola, Alessandro, Georgi Kocharkov, Almuth Scholl, and Anna-Mariia Tkhir. 2021. “The Aggregate Consequences of Tax Evasion.” *Review of Economic Dynamics* 40:198–227.
- Diamond, Peter A. 1967. “The Role of a Stock Market in a General Equilibrium Model with Technological Uncertainty.” *The American Economic Review* 57 (4):759–776.
- Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri. 2020. “Heterogeneity and Persistence in Returns to Wealth.” *Econometrica* 88 (1):115–170.
- French, Eric and John Bailey Jones. 2011. “The Effects of Health Insurance and Self-Insurance on Retirement Behavior.” *Econometrica* 79 (3):693–732.
- Gallipoli, Giovanni, Hamish Low, and Aruni Mitra. 2020. “Consumption and Income Inequality Across Generations.” .
- Gourinchas, Pierre-Olivier and Jonathan A. Parker. 2002. “Consumption Over the Life Cycle.” *Econometrica* 70 (1):47–89.
- Güvener, Fatih, Gueorgui Kambourov, Burhan Kuruscu, and Sergio Ocampo. 2022a. “Taxing Wealth and Capital Income when Returns are Heterogeneous.” .

- Guvenen, Fatih, Gueorgui Kambourov, Burhanettin Kuruscu, Sergio Ocampo-Diaz, and Daphne Chen. 2022b. “Use It or Lose It: Efficiency Gains from Wealth Taxation.” Working Paper.
- Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song. 2021. “What Do Data on Millions of U.S. Workers Reveal About Lifecycle Earnings Dynamics?” *Econometrica* 89 (5):2303–2339.
- Guvenen, Fatih, Burhanettin Kuruscu, and Serdar Ozkan. 2014. “Taxation of Human Capital and Wage Inequality: A Cross-Country Analysis.” *Review of Economic Studies* 81 (2):818–850.
- Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2014. “Consumption and Labor Supply with Partial Insurance: An Analytical Framework.” *American Economic Review* 104 (7):2075–2126.
- . 2017. “Optimal Tax Progressivity: An Analytical Framework.” *Quarterly Journal of Economics* 132 (4):1693–1754.
- Holter, Hans A, Dirk Krueger, and Serhiy Stepanchuk. 2019. “How Do Tax Progressivity and Household Heterogeneity Affect Laffer Curves?” *Quantitative Economics* 10 (4):1317–1356.
- Jordà, Òscar, Katharina Knoll, Dmitry Kuvshinov, Moritz Schularick, and Alan M. Taylor. 2019. “The Rate of Return on Everything, 1870–2015.” *Quarterly Journal of Economics* 134 (3):1225–1298.
- Kaplan, Greg and Giovanni L. Violante. 2014. “A Model of the Consumption Response to Fiscal Stimulus Payments.” *Econometrica* 82 (4):1199–1239.
- Kopecky, Karen A. and Richard M.H. Suen. 2010. “Finite State Markov-Chain Approximations to Highly Persistent Processes.” *Review of Economic Dynamics* 13 (3):701–714.
- Krueger, Dirk, Alexander Ludwig, and Sergio Villalvazo. 2021. “Optimal taxes on capital in the OLG model with uninsurable idiosyncratic income risk.” *Journal of Public Economics* 201:104491.
- Laczó, Sarolta and Raffaele Rossi. 2020. “Time-Consistent Consumption Taxation.” *Journal of Monetary Economics* 114:194–220.
- Leung, David and Markus Poschke. 2022. “Progressive Consumption Tax Reforms.” Mimeo.

- Lucas, Robert E. 1987. *Models of Business Cycles*, vol. 26. Basil Blackwell Oxford.
- McDaniel, Cara. 2007. “Average Tax Rates on Consumption, Investment, Labor and Capital in the OECD 1950-2003.” Working Paper.
- McFadden, Daniel. 1978. “Modelling the Choice of Residential Location.” In *Spatial Interaction Theory and Planning Models*, edited by A. Karlqvist, L. Lundqvist, F. Snickars, and J. Weibull. North Holland, 75–96.
- Moll, Benjamin. 2014. “Productivity Losses from Financial Frictions: Can Self-Financing Undo Capital Misallocation?” *American Economic Review* 104 (10):3186–3221.
- Nishiyama, Shinichi and Kent Smetters. 2005. “Consumption Taxes and Economic Efficiency with Idiosyncratic Wage Shocks.” *Journal of Political Economy* 113 (5):1088–1115.
- OECD. 2020. *Consumption Tax Trends 2020*. URL <https://www.oecd-ilibrary.org/content/publication/152def2d-en>.
- Peterman, William B and Erick Sager. 2022. “Optimal Public Debt with Life Cycle Motives.” *American Economic Journal: Macroeconomics* 14 (4):404–37.
- Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman. 2018. “Distributional National Accounts: Methods and Estimates for the United States.” *Quarterly Journal of Economics* 133 (2):553–609.
- Saez, Emmanuel and Gabriel Zucman. 2016. “Wealth Inequality in the United States Since 1913: Evidence from Capitalized Income Tax Data.” *The Quarterly Journal of Economics* 131 (2):519–578.
- Smith, Matthew, Owen M Zidar, and Eric Zwick. 2021. “Top Wealth in America: New Estimates and Implications for Taxing the Rich.” Tech. rep., National Bureau of Economic Research.
- Summers, Lawrence H. 1981. “Capital Taxation and Accumulation in a Life Cycle Growth Model.” *American Economic Review* 71 (4):533–544.
- Tauchen, George. 1986. “Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions.” *Economics Letters* 20 (2):177–181.
- Xavier, Inês. 2020. “Wealth Inequality in the US: The Role of Heterogeneous Returns.” Working Paper.

Appendix

A Definition of Equilibrium

We focus on a stationary equilibrium, in which capital, labor, transfers, and government consumption are all constant in per-capita terms. Let $\psi_{i,j}(a, z_h, z_r)$ denote the distribution of agents with type i and age j , over assets a and idiosyncratic shocks (z_h, z_r) .

Definition 1. *The stationary recursive equilibrium consists of*

- (i) *the value function, $V_{i,j}(a, z_h, z_r)$;*
- (ii) *the policy functions, $c_{i,j}(a, z_h, z_r)$, $a'_{i,j}(a, z_h, z_r)$, $h_{i,j}(a, z_h, z_r)$;*
- (iii) *the entrepreneurial profit function $\pi(a, x)$ and associated capital demand $k(a, x)$;*
- (iv) *the prices (w, p, r) ;*
- (v) *the per-capita stocks of capital K , intermediate good Q , labor L , government spending G ;*
- (vi) *the the social security benefit level χ ;*
- (vii) *the per-capita benefit levels b_i , lump-sum transfers $T_{b,i}$ and labor L_i for types $i = 1, \dots, I$; and*
- (viii) *distributions (μ_1, \dots, μ_J) , $(\psi_{i,1}, \dots, \psi_{i,J})$ for $i = 1, \dots, I$*

such that the following conditions hold.

1. *The value function $V_{i,j}(a, z_h, z_r)$ solves the Bellman equation in (13) and $c_{i,j}(a, z_h, z_r)$, $a'_{i,j}(a, z_h, z_r)$, $h_{i,j}(a, z_h, z_r)$ are the associated policy functions.*
2. *Household profits $\pi(a, x)$ solve (6) and capital demand $k(a, x)$ is given by (7).*
3. *The final goods producer maximizes its profits, requiring that $F_Q(Q, L) = p$ and $F_L(Q, L) = w$.*

4. Markets clear:

$$\begin{aligned}
\sum_{i=1}^I \pi_i \sum_{j=1}^J \mu_j \int [c_{i,j}(a, z_h, z_r) + a'_{i,j}(a, z_h, z_r)] d\psi_{i,j} + G &= F(Q, L) + (1 - \delta)K \\
\sum_{i=1}^I \pi_i \sum_{j=1}^J \mu_j \int z_r k(a, z_r) d\psi_{i,j} &= Q \\
\sum_{i=1}^I \pi_i \sum_{j=1}^J \mu_j \int a d\psi_{i,j} &= K \\
\sum_{i=1}^I \pi_i \sum_{j=1}^J \mu_j \int [k(a, x_j(z_r)) - a] d\psi_{i,j} &= 0 \\
\sum_{j=1}^J \mu_j \int e_{i,j}(z_h) h_{i,j}(a, z_h, z_r) d\psi_{i,j} &= L_i \\
\sum_{i=1}^I \pi_i L_i &= L.
\end{aligned}$$

5. The distribution of agents across age groups, μ_1, \dots, μ_J , satisfies

$$\mu_{j+1} = \frac{s_{j+1}\mu_j}{1+n} \text{ for } j = 1, \dots, J-1$$

where μ_1 is normalized so that $\sum_{t=1}^J \mu_j = 1$.

6. The distributions of agents within each age group j and type i , $\psi_{i,1}, \dots, \psi_{i,J}$, for $i = 1, \dots, I$, are consistent with individual behavior. That is, the law of motion for $\psi_{i,j}$ is

$$\psi_{i,j+1}(a', z'_h, z'_r) = \int f(z'_h|z_h) f(z'_r|z_r) \mathbb{1} \{a'_{i,j}(a, z_h, z_r) = a'\} d\psi_{i,j}(a, z_h, z_r)$$

where $f(z'_h|z_h)$ and $f(z'_r|z_r)$ are the conditional probabilities for the household transitioning to z'_h and z'_r given that its current shocks are z_h and z_r , respectively.

Furthermore, in the initial distribution $\psi_{i,1}(a, z_h, z_r)$ for each type $i \in \{1, \dots, I\}$, all age-1 agents are born with no assets (i.e., $a = 0$), the initial labor productivity shock, $\log z_h$, is zero and the initial $\log z_r$ is drawn from $N(0, \sigma_{\varepsilon_r}^2 / (1 - \rho_r^2))$.

7. The government budget constraint is satisfied

$$G = T_y \equiv \sum_{i=1}^I \pi_i \sum_{j=1}^J \mu_j [T_l(y_{l,i,j}(a, z_h, z_r)) + \tau_c c_{i,j}(a, z_h, z_r) + \tau_a + \tau_k r_j^a(z_r) a] d\psi_{i,j}$$

where taxable labor income is

$$y_{l,i,j}(a, z_h, z_r) = we_{i,j}(z_h)h_{i,j}(a, z_h, z_r) - \frac{1}{2}\tau_{ss} \min (we_{i,j}(z_h)h_{i,j}(a, z_h, z_r), \bar{y}).$$

8. Social security benefits equal social security taxes:

$$\tau_{ss} \sum_{i=1}^I \pi_i \sum_{j=1}^J \int \min (we_{i,j}(z_h)h_{i,j}(a, z_h, z_r), \bar{y}) d\psi_{i,j} = \sum_{i=1}^I \pi_i b_i \left(\sum_{j=R}^J \mu_j \right).$$

9. The social security benefit levels are $\bar{b}_i = \chi \Phi (\min (wL_i, \bar{y}))$.

10. Lump-sum transfers T_b are consistent with individual behavior,

$$T_{b,i} = \frac{1}{1+n} \sum_{j=1}^J \mu_j (1 - s_{j+1}) \int a'_{i,j}(a, z_h, z_r) (1+r) d\psi_{i,j}.$$

B Numerical Solution Technique

The numerical solution technique is standard. First, we describe the discrete approximations we make for the idiosyncratic shocks and the fixed levels of innate ability. Second, we describe how we solve for the stationary equilibrium.

Discrete Approximations First, we discretize the AR(1) process for the entrepreneurial shock z_r using the Rouwenhorst method (see [Kopecky and Suen, 2010](#)). While the labor ability shock follows a random walk, we can still discretize z_h as households work for a finite number of periods. That is, we construct the grid for $\ln z_h$ as a linearly spaced vector between $-3\sigma_{\varepsilon_h} \sqrt{R-1}$ and $+3\sigma_{\varepsilon_h} \sqrt{R-1}$. Then, we construct a probability transition matrix following the approach of [Tauchen \(1986\)](#).

Second, we discretize the fixed levels of labor ability. That is, given the standard deviation of innate ability σ_e , we set $\{\bar{e}_i\}_{i=1}^I$ as I individual points, linearly spaced between $-3\sigma_e$ and $+3\sigma_e$. Second, assuming innate labor ability is normally distributed with mean zero and variance σ_e^2 , we construct the individual type probabilities $\{\pi_i\}_{i=1}^I$ using the approximation method of [Tauchen \(1986\)](#).

Approximating the Entrepreneurial Decision In order to enable convergence when solving for the stationary equilibrium, we approximate the household's entrepreneurial decision by introducing two additional i.i.d. shocks, $(\varepsilon_e, \varepsilon_b)$. These two shocks are independent

of each other and are assumed to be drawn from a generalized extreme value (type-I) distribution. That is, we assume the following cumulative distribution functions:

$$F(\varepsilon_e) = \exp \left[- \exp \left(- \frac{\varepsilon_e - \mu_\varepsilon}{\sigma_\varepsilon} \right) \right]$$

$$F(\varepsilon_b) = \exp \left[- \exp \left(- \frac{\varepsilon_b - \mu_\varepsilon}{\sigma_\varepsilon} \right) \right]$$

The parameter μ_ε is a location parameter and σ_ε is a scale parameter which controls the variance of the shocks. We set σ_ε to a small value (0.001) and set $\mu_\varepsilon = -\sigma_\varepsilon\gamma_\varepsilon$ where $\gamma_\varepsilon \approx 0.577$ is Euler's constant. This ensures that the mean value of the shocks is zero. These shocks, while small, will randomize the decision over whether to be an entrepreneur, mainly for those households who are close to being indifferent between being an entrepreneur or not.

We introduce these shocks into the model as follows. Consider a household, with entrepreneurial productivity x , choosing whether to be an entrepreneur or not. We assume the return on assets when an entrepreneur (r_e) and the return when not an entrepreneur (r_b) are given by:

$$r_e = \bar{r}_e(x) + \varepsilon_e \quad \text{where} \quad \bar{r}_e(x) \equiv r + \lambda(px - r - \delta)$$

$$r_b = r + \varepsilon_b.$$

With extreme value shocks, the entrepreneurial decision is now random. If the household becomes an entrepreneur, it invests $k = \lambda a$ in its backyard technology. The probability that a household will choose to be an entrepreneur depends on its entrepreneurial productivity and is given by:

$$p_e(x) = \frac{\exp(\bar{r}_e(x)/\sigma_\varepsilon)}{\exp(\bar{r}_e(x)/\sigma_\varepsilon) + \exp(r/\sigma_\varepsilon)}. \quad (\text{B.1})$$

Therefore, the return on assets for an age- j household with entrepreneurial shock z_r is then

$$r_j^a(z_r) = \mu_\varepsilon + \sigma_\varepsilon\gamma_\varepsilon + \sigma_\varepsilon \ln [\exp(\bar{r}_e(x_j(z_r))/\sigma_\varepsilon) + \exp(r/\sigma_\varepsilon)] \quad (\text{B.2})$$

where $\gamma_\varepsilon \approx 0.577$ is Euler's constant and $\mu_\varepsilon + \sigma_\varepsilon\gamma_\varepsilon = 0$ by construction.²⁵

Solving for the Stationary Equilibrium We solve for the stationary equilibrium using an iterative Gauss-Jacobi algorithm. That is, given a guess for the variables $\mathbf{y} = (\bar{x}, \chi, Q, \{T_{b,i}\}_{i=1}^I, \{L_i\}_{i=1}^I)$, we solve for the household's value functions and policy rules, simulate the economy and construct a new guess \mathbf{y}' . We repeat the process until \mathbf{y} con-

²⁵To derive Equations (B.1) and (B.2), we use Theorem 1 in [McFadden \(1978\)](#).

verges. More specifically, the algorithm proceeds as follows.

1. We start with a guess for the vector $\mathbf{y} = (\bar{x}, \chi, Q, \{T_{b,i}\}_{i=1}^I, \{L_i\}_{i=1}^I)$.
2. Given $\{L_i\}_{i=1}^I$, we compute aggregate labor $L = \sum_{i=1}^I \pi_i L_i$.
3. Given Q and L , we determine prices $p = F_Q(Q, L)$ and $w = F_L(Q, L)$.
4. Given χ , w and L_i , we determine the social security benefit $b_i = \chi \Phi(\min(wL_i, \bar{y}))$.
5. Given \bar{x} and p , we compute the borrowing rate $r = p\bar{x} - \delta$.
6. Given w , r , p , b_i , $T_{b,i}$, we solve for the policy functions $a'_{i,j}(a, z_h, z_r)$, $h_{i,j}(a, z_h, z_r)$ for $i = 1, \dots, I$, $j = 1, \dots, J$ by iterating on the Bellman equation defined in (13). We use Schumaker interpolation to interpolate the value function over assets. To determine the optimal choices of hours and savings, we use grid search followed by the BOBYQA local minimization algorithm. Consumption is then determined by the household's budget constraint.
7. We calculate the distributions $\psi_{i,j}$ for $i = 1, \dots, I$ and $j = 1, \dots, J$ using Monte Carlo simulation.
8. Given $\psi_{i,j}$ and \mathbf{y} , we construct an updated guess $\hat{\mathbf{y}} = (\hat{\bar{x}}, \hat{\chi}, \hat{Q}, \{\hat{T}_{b,i}\}_{i=1}^I, \{\hat{L}_i\}_{i=1}^I)$ as follows:
 - (a) We update the guess for the social security benefit parameter using the social security budget constraint:

$$\hat{\chi} = \left[\sum_{i=1}^I \pi_i \sum_{j=1}^J \mu_j \int \tau_{ss} \min(w e_{i,j}(z_h) h_{i,j}(a, z_h, z_r), \bar{y}) d\psi_{i,j} \right] \times \left[\left(\sum_{i=1}^I \pi_i \Phi(\min(wL_i, \bar{y})) \right) \left(\sum_{j=1}^J \mu_j \right) \right]^{-1}.$$

- (b) We update the guess for quality-adjusted capital:

$$\hat{Q} = \sum_{i=1}^I \pi_i \sum_{j=1}^J \mu_j \int x_j(z_r) p_e(x_j(z_r)) \lambda a d\psi_{i,j}$$

where $p_e(x_j(z_r))$ is the probability a household with entrepreneurial productivity $x_j(z_r)$ will choose to be an entrepreneur. Note that $p_e(x)$ is defined in Equation (B.1).

(c) We update the guess for aggregate per-capita labor of type i :

$$\hat{L}_i = \sum_{j=1}^J \mu_j \int e_{i,j}(z_h) h_{i,j}(a, z_h, z_r) d\psi_{i,j} \text{ for } i = 1, \dots, I$$

(d) We update the guess for the lump-sum transfer to type i :

$$\hat{T}_{b,i} = \frac{1}{1+n} \left[\sum_{j=1}^J \mu_j (1 - s_{j+1}) \int a'_{i,j}(a, z_h, z_r) (1+r) d\psi_{i,j} \right]$$

(e) We update the guess for the entrepreneurial cutoff:

$$\hat{x} = \bar{x} \left[1 + \eta \left(\frac{\hat{K}_e - \hat{K}}{\hat{K}} \right) \right]$$

where

$$\begin{aligned} \hat{K}_e &= \sum_{i=1}^I \pi_i \sum_{j=1}^J \mu_j \int p_e(x_j(z_r)) \lambda a d\psi_{i,j} \\ \hat{K} &= \sum_{i=1}^I \pi_i \sum_{j=1}^J \mu_j \int a d\psi_{i,j}. \end{aligned}$$

Note that \hat{K}_e is the aggregate capital demanded by entrepreneurs, while \hat{K} is the aggregate wealth supplied by households. In equilibrium, we require $\hat{K}_e = \hat{K}$. Therefore, we internally use a parameter $\eta = 0.29$ in the numerical algorithm to update the guess for \bar{x} . When there is excess demand for capital from entrepreneurs, the algorithm will tend to increase \bar{x} and thus increase the borrowing rate. When there is excess supply, the opposite will occur.

9. If \mathbf{y} is sufficiently close to $\hat{\mathbf{y}}$, the algorithm stops. Otherwise, we update \mathbf{y} , and return to step 2. To ensure stability, we use a convex combination:

$$\mathbf{y}' = \omega \hat{\mathbf{y}} + (1 - \omega) \mathbf{y}$$

where $\omega \in (0, 1)$. Generally, we set $\omega = 0.25$ but we utilized a procedure to decrease ω in increments of 0.05 if the algorithm was not successfully converging (i.e., the guesses were bouncing back and forth around the solution).

C Additional Estimates

C.1 Survival Probabilities

The survival probabilities were obtained from the 2018 Period Life Tables from United States Mortality Database (see Table C.1). We utilized survival probabilities for both genders across the entire United States. Since the maximum age is $J = 85$ in the model (which corresponds to age 105 in real life), we impose that $s_{J+1} = 0$.

Table C.1 – Survival Probabilities

Age (j)	s_{j+1}	Age (j)	s_{j+1}	Age (j)	s_{j+1}
1	0.9991	31	0.9957	61	0.9495
2	0.9991	32	0.9953	62	0.9435
3	0.9990	33	0.9949	63	0.9371
4	0.9990	34	0.9943	64	0.9326
5	0.9989	35	0.9938	65	0.9261
6	0.9989	36	0.9932	66	0.9175
7	0.9988	37	0.9927	67	0.9077
8	0.9988	38	0.9922	68	0.8952
9	0.9987	39	0.9916	69	0.8837
10	0.9987	40	0.9909	70	0.8699
11	0.9987	41	0.9901	71	0.8547
12	0.9986	42	0.9895	72	0.8408
13	0.9985	43	0.9887	73	0.8230
14	0.9985	44	0.9881	74	0.8054
15	0.9984	45	0.9872	75	0.7919
16	0.9983	46	0.9863	76	0.7732
17	0.9983	47	0.9855	77	0.7537
18	0.9982	48	0.9846	78	0.7335
19	0.9981	49	0.9832	79	0.7128
20	0.9980	50	0.9819	80	0.6919
21	0.9979	51	0.9790	81	0.6705
22	0.9979	52	0.9794	82	0.6488
23	0.9977	53	0.9763	83	0.6275
24	0.9976	54	0.9739	84	0.6076
25	0.9973	55	0.9702	85	0
26	0.9971	56	0.9689		
27	0.9969	57	0.9656		
28	0.9967	58	0.9627		
29	0.9964	59	0.9585		
30	0.9961	60	0.9538		

C.2 Labor Tax Function Estimation

We employ a hybrid approach to estimate the two parameters of our labor tax function, given in Equation (10). We directly estimate the progressivity of the labor income tax schedule (τ_l) using the 2019 Survey of Consumer Finances (SCF), as we describe below. To estimate the level parameter, λ_l , we include it in our SMM estimation. To identify this parameter, we target the average labor income tax rate from the 2019 SCF.

To estimate the progressivity parameter τ_l , we regress log post-tax labor income on log pre-tax labor income. The slope of this regression is $1 - \tau_l$. To determine the federal income taxes that can be interpreted as taxes on labor income, we use NBER's TAXSIM program. However, when computing federal tax liabilities, we assume households have no interest income, dividends, capital gains or other property income. In our estimation, we restrict the sample households whose age is between 21 and 64, which corresponds to the working ages in our model. Then, we regress log post-tax labor income on log pre-tax income. With this estimation, we obtain $\tau_l = 0.1995$.

D Welfare Decomposition

We compute the welfare gain by computing the percentage increase in consumption required in the initial economy, every period, denoted by CEV , that makes a household indifferent between the consumption and hours path of the old economy (c_0, h_0) and the new economy (c_*, h_*).

$$W(c_*, h_*) = W(c_0(1 + CEV), h_0)$$

Next, we decompose the overall change into effects due to consumption and labor supply (leisure): $1 + CEV = (1 + CEV_c) * (1 + CEV_h)$:

$$W(c_*, h_0) = W(c_0(1 + CEV_c), h_0)$$

$$W(c_*, h_*) = W(c_*(1 + CEV_h), h_0)$$

Viewed another way:

$$\underbrace{W(c_0, h_0) \xrightarrow{CEV_c} W(c_*, h_0) \xrightarrow{CEV_h} W(c_*, h_*)}_{CEV}$$

Then, we decompose the consumption effect into a level effect CEV_{cL} and a distribution effect CEV_{cD} . Define $\hat{c}_0 = (C_*/C_0)c_0$ as the consumption allocation resulting from scaling

the allocation c_0 by the change in aggregate consumption C_*/C_0 .

$$\begin{aligned} W(\hat{c}_0, h_0) &= W(c_0(1 + CEV_{cL}), h_0) \\ W(c_*, h_0) &= W(\hat{c}_0(1 + CEV_{cD}), h_0) \end{aligned}$$

Similarly, we can decompose the labor supply effect into a level effect CEV_{hL} and a distribution effect CEV_{hD} . Define $\hat{h}_0 = (H_*/H_0)h_0$ as the hours allocation resulting from scaling the allocation h_0 by the change in aggregate hours H_*/H_0 .

$$\begin{aligned} W(c_*, \hat{h}_0) &= W(c_*(1 + CEV_{hL}), h_0) \\ W(c_*, h_*) &= W(c_*(1 + CEV_{hD}), \hat{h}_0) \end{aligned}$$

Viewed in total:

$$\underbrace{W(c_0, h_0) \xrightarrow{CEV_{cL}} W(\hat{c}_0, h_0) \xrightarrow{CEV_{cD}} W(c_*, h_0)}_{CEV_c} \underbrace{\xrightarrow{CEV_{hL}} W(c_*, \hat{h}_0) \xrightarrow{CEV_{hD}} W(c_*, h_*)}_{CEV_h} \xrightarrow{CEV} W(c_*, h_*)$$

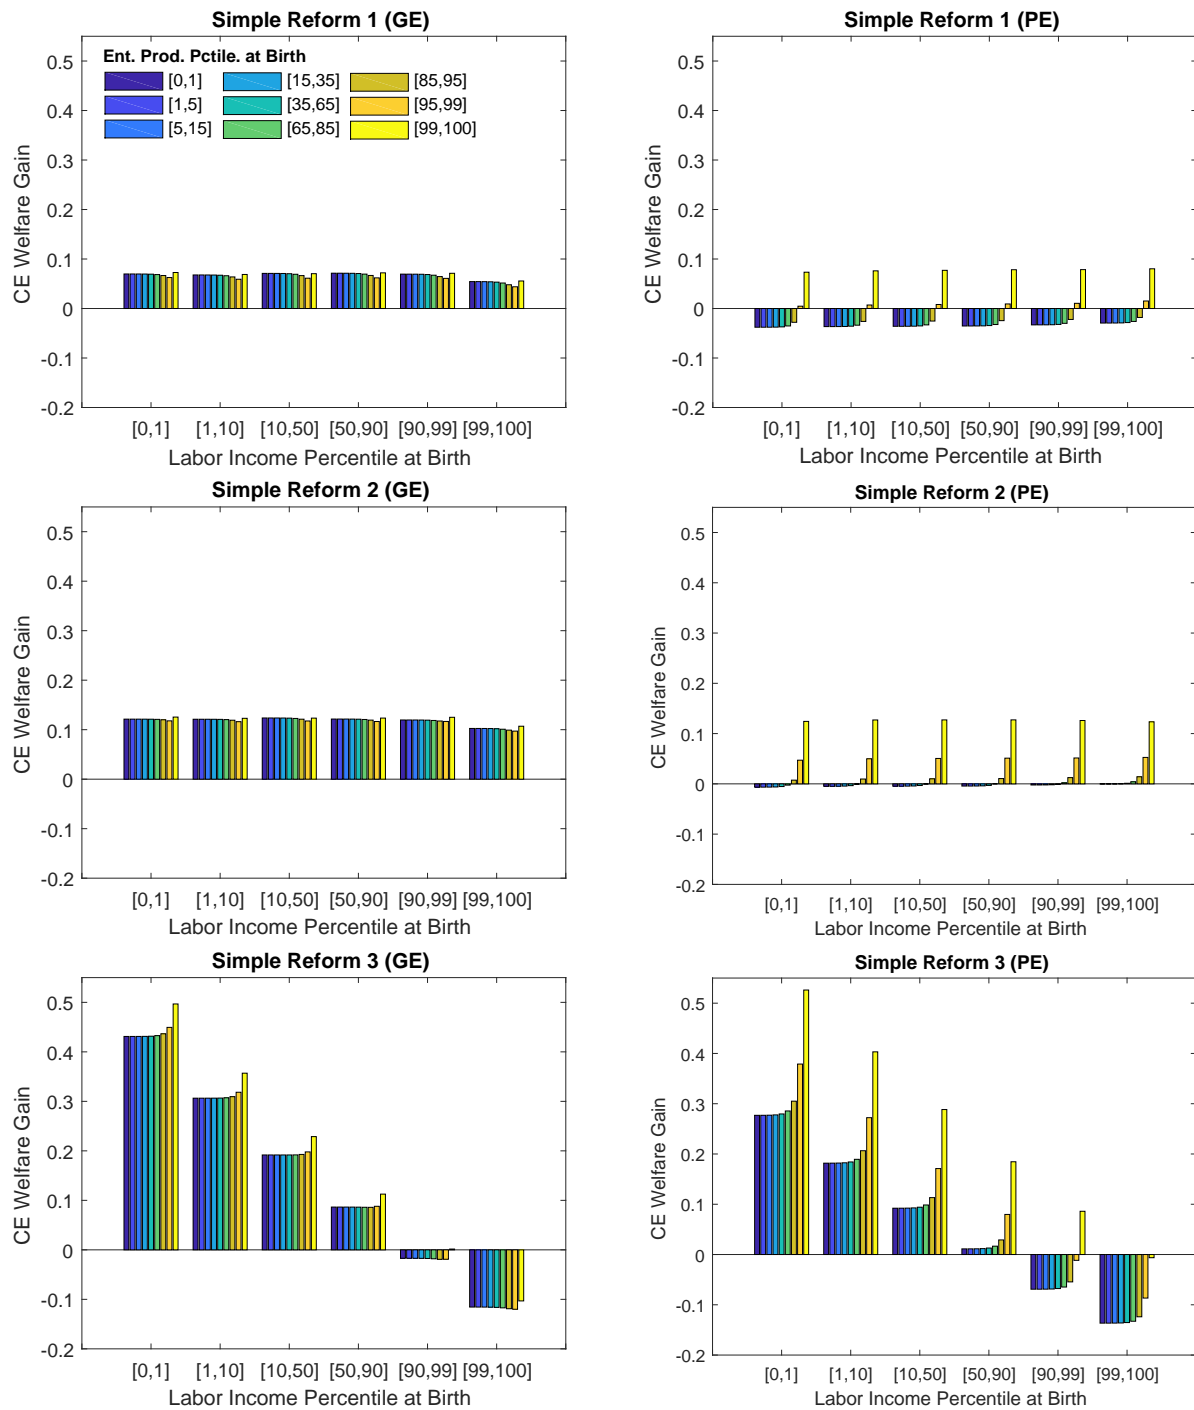
E Fiscal Reforms in Partial Equilibrium

Here we analyze the same partial reform as in the benchmark exercise, but in partial equilibrium (or small open economy) – i.e., we keep prices fixed. Results from this exercise are presented in Figure E.1. Notice that in partial equilibrium, there is a lot of heterogeneity in the welfare gains across individuals. Especially for the first two simple reforms, general equilibrium effects are crucial for explaining why the welfare gains are so similar across all individuals. Absent any general equilibrium effects, it would only be individuals with the highest entrepreneurial ability who would benefit from the first two simple reforms. For the final reform, general equilibrium effects reduce the gains to high-entrepreneurial ability households and increase the gains to low-entrepreneurial ability households.

F Robustness Checks with Alternative Models

Here we conduct extensive robustness and show the optimality of consumption taxation by changing various critical features of the baseline economy. The take home from these experiments is that our policy proposal is robust to a large number of changes that have been explored by the literature. Given that parameters are kept at their baseline values, each model will, by construction, have different distributional and aggregate properties. These

Figure E.1 – Welfare Gains of Simple Reforms in Partial Equilibrium



Note: This figure reports the welfare gains from simple reforms in general equilibrium (left panels) and in partial equilibrium (right panels).

will be briefly discussed on a case-by-case basis.²⁶ Results from this exercise are reported in Table F.1.

The Case with No Return Heterogeneity. We start our exercise by analyzing a crucial feature in our environment, namely stochastic return to savings via fluctuations in entrepreneurial ability and financial constraints. As previously discussed in this paper and comprehensively summarized in [Benhabib and Bisin \(2018\)](#), return heterogeneity is a crucial (and empirically consistent) mechanism for replicating the thickness of the right tail of the income and wealth distributions. It also creates a scope for shifting the burden of taxation from capital income to consumption or in absence of the latter to wealth. A natural question is therefore what happens to the tax reform put forward in this paper in a model all equal to the baseline one, but where return to wealth is constant across individuals.

Specifically, we assume $\kappa_1 = \kappa_2 = 0$ (i.e., there is no age profile to entrepreneurial productivity) and $\sigma_{\varepsilon r} = 0$ (no idiosyncratic uncertainty in entrepreneurial productivity). As a result, all individuals have the same level of entrepreneurial productivity, x , but we normalize this level so that it is equal to Q/K from our benchmark economy (as a result, aggregate TFP will be identical in the two economies). With these modifications, this economy is equivalent to the conventional framework in which all households earn the same rate of return (r) on their savings.

Interestingly, optimal policy looks remarkably similar to the benchmark case. The planner still finds it optimal to raise all revenues with consumption taxation and provide social insurance via the same progressive wage tax schedule. A few considerations are in order. First, these results indicate that the optimality of increasing progressivity in the wage tax schedule relative to the status quo social is mainly driven by the combination of the tax structure and earning risk, which are kept here as in our benchmark exercise. Second, absent of return heterogeneity, optimal policy has a detrimental effect on output (-3%) and a stronger negative effect on private consumption (-10.3%). From a welfare point of view, most of the gains come from redistribution (19%), while the absence of the efficiency gains from the reallocation of capital shrinks the welfare on the level effect (-2.9%).²⁷

These results point out that while there are quantitative effects in shutting down the inefficiency of capital allocation, this is not the main driver behind our results, and taxing consumption remains very appealing. This is important because, as reported in [Boar and](#)

²⁶One potential alternative would have been to re-estimate the model for each scenario. However, for many cases, we would not be able to match the distributional and aggregate properties of the data.

²⁷We also an alternative (and less extreme) experiment in which we shut down only the age component of the entrepreneurial abilities and results remain, for all practical purposes as in the benchmark, so we decided to not including it for sake of brevity.

Midrigan (2022), different modeling assumptions behind the importance of capital misallocation might twist the relative merit of taxing wealth vis-a-vis taxing capital income.

The Case of No Transitory Earning Risk. For this exercise, we set $\sigma_{\varepsilon h} = 0$, so that agents do not face any transitory risk in their labor productivity. Nevertheless, labor productivity will still depend on a deterministic age profile and an innate household-specific fixed effect. In this case, the need of social insurance is greatly reduced, as the role of missing markets is less detrimental on welfare. Interestingly, the planner still finds it beneficial to provide redistribution through a progressive wage tax schedule, although the optimal progressivity is greatly reduced relative to the benchmark exercise, to 0.28. These results indicate that both permanent and transitory earning heterogeneity call for redistribution. The qualitative effects of optimal policy are comparable *ceteris paribus* to the benchmark case. Given the reduced level of risk in the economy, the gains from the optimal tax schedule are greatly reduced to 8.6%, and mostly due to level effects. Surprisingly, these results indicate that a large chunk of the welfare benefits of the optimal policy are due to redistribution over earning risk. In this sense, it is crucial the role of consumption taxation in increasing the optimal progressivity of the wage tax schedule.

The Case with Decreasing Return to Scale. Our micro-foundation of return heterogeneity through backyard technology relies on constant return to scale. A natural question is to understand what happens under the alternative assumption of decreasing return to scale (DRS). Given we want to maintain entrepreneurial decision that varies with age and abilities, we need to have a fixed cost in setting up business. This avoids having everyone being an entrepreneur, like in Guvenen et al. (2022b), or a once-for-life entrepreneurial decision, like Cagetti and De Nardi (2006), and makes the model as close as possible to our benchmark, thus isolating in a transparent manner the role of return to scale.

Specifically, we assume entrepreneurial households produce the intermediate capital service using a decreasing-returns-to-scale technology:

$$q = x_j(z_r)k^\nu - c_f$$

where $\nu \in (0, 1)$ is the scale parameter and c_f is a fixed cost the household incurs if it chooses to be an entrepreneur. Entrepreneurial profits (if the household chooses to become an entrepreneur) are now

$$\pi(a, x) = \max_{0 \leq k \leq \lambda a} \{p(xk^\nu - c_f) - (r + \delta)k\}.$$

Because of the fixed cost, $\pi(a, x)$ will be negative for low-productivity households, and thus these households will choose not to enter entrepreneurship. Household returns on wealth are now

$$r_j(z_r, a) = \begin{cases} r + \frac{1}{a} \max(\pi(a, x_j(z_r)), 0) & \text{if } a > 0 \\ r & \text{if } a = 0. \end{cases}$$

The shape of the optimal policy is qualitatively similar to the benchmark. However, relative to the benchmark with CRS, the efficiency gains from reallocation of capital are reduced with decreasing returns, and the overall increase on TFP is lower. This is because with DRS there is a mechanical dampening effect on aggregate productivity in reallocating capital towards the most productive entrepreneurs. Furthermore, contrary to the benchmark case with CRS, output decreases (-3.3%) in the optimal policy. The lower efficiency gains from this model reduce the overall “level” welfare gains to 0.3%, while the distributional gains from the policy are magnified, relative to the benchmark case. This is because the distortionary effects of higher progressivity is lower, as the efficient allocation places less weight on redistributing capital to the top ability entrepreneurs.

The Case with Intergenerational Links and Joint Distributions of Abilities. We now present a similar model in which we allow for (i) intergenerational links in abilities; and (ii) joint distribution between labor and entrepreneurial abilities. We present these two ingredients jointly, but we experimented with each feature in isolation, and results do not change.

For this experiment, we assume explicit links between generations, where the wealth of dying households is transmitted to its heirs as a bequest. For this purpose, we add a fixed effect to entrepreneurial productivity, \bar{x}_i , so that the latter is now given by:

$$\log x_{i,j}(z_r) = \bar{x}_i + \kappa_1 \left(j - \frac{1}{J} \sum_{j'=1}^J j' \right) + \kappa_2 \left(j^2 - \frac{1}{J} \sum_{j'=1}^J (j')^2 \right) + \log z_r.$$

As a result, there is now fixed heterogeneity at birth across both labor productivity (\bar{e}_i) and entrepreneurial productivity (\bar{x}_i), where i indexes the agent’s type i .

We assume the fixed-effect components of both labor productivity and entrepreneurial productivity are linked between parents and children via an autoregressive process:

$$\begin{bmatrix} \bar{e}^{child} \\ \bar{x}^{child} \end{bmatrix} = \begin{bmatrix} \rho_{\bar{e}} & 0 \\ 0 & \rho_{\bar{x}} \end{bmatrix} \begin{bmatrix} \bar{e}^{parent} \\ \bar{x}^{parent} \end{bmatrix} + \varepsilon, \quad \varepsilon \sim N(0, \Sigma) \quad \text{with } \Sigma = \begin{bmatrix} \sigma_{\bar{e}}^2 & \rho_{\bar{e}, \bar{x}} \sigma_{\bar{e}} \sigma_{\bar{x}} \\ \rho_{\bar{e}, \bar{x}} \sigma_{\bar{e}} \sigma_{\bar{x}} & \sigma_{\bar{x}}^2 \end{bmatrix}$$

The unconditional standard deviation of \bar{e} is then equal to $\sigma_{\bar{e}} / \sqrt{1 - \rho_{\bar{e}}^2}$, which we set to

σ_e from our benchmark economy. We pick $\rho_{\bar{e}} = 0.25$ and $\rho_{\bar{x}} = 0.25$, which control the persistence of labor and entrepreneurial productivity, respectively, across generations. We set $\rho_{\bar{e},\bar{x}} = 0.5$, which controls the correlation between entrepreneurial and labor ability. We fixed the intergenerational link in labor ability from the recent evidence provided by [Gallipoli et al. \(2020\)](#). We are not aware of any reliable estimates of the intergenerational link in entrepreneurial abilities, so we fixed it to the same level of labor abilities. Similarly, we are not aware of any estimate on the joint distribution of abilities, so we agnostically choose 0.5. And finally, as in our benchmark economy, we adapted the approach of [Tauchen \(1986\)](#) to discretize this process to generate $i \in \{1, \dots, I\}$ types which now index both labor and entrepreneurial ability.

Interestingly, optimal policy looks remarkably similar to the benchmark case. This said, there are some small quantitative differences. First of all, optimal progressivity in wage tax schedule increases marginally, as intergenerational links increase the optimal level of redistribution, e.g., consumption gini decreases more than in the benchmark case (-12.2% vs. 10.7%). As such, the negative level effect on the consumption component of welfare, CEV_{cL} (not reported in the table), is stronger than in the benchmark (-4.5% vs. -2.9%). This results into a lower overall level effect on the overall welfare (2.7% vs. 4%), while the gains from redistribution are larger (14.4% vs 13%).

The Case with Fixed Average Wage Taxes. One natural question one may raise is whether our results on optimal progressivity depend on a very low average labor income tax. To address this concern, here conduct an alternative experiment, whereas we skip the second partial reform [4.1](#). Instead, we optimally tilt the progressivity of wage taxation leaving the average wage tax at the status-quo of 13.4%. As in all other cases, any fiscal requirement coming from the change in progressivity is compensated by increasing consumption taxation. It is important to stress we conduct this step after having eliminated capital income taxation.

Somewhat remarkably, the optimal progressivity of the labor income tax schedule turns out like in the benchmark exercise, at 0.35. This means that allowing consumption taxation to adjust creates the incentive for the planner to increase social insurance in the labor market. This is a clear manifestation of consumption taxes as efficient means of fiscal revenue generator. Not surprisingly, the welfare gains from this policy are smaller than in the benchmark. This is because the planner is losing from level and distributional gains by keeping average wage tax at the suboptimal high level.

Table F.1 – Optimal Constrained Policy: Robustness

	Bench.	No Ret. Het.	No Wage Trans. Risk	DRS	Inter. Links + Corr. Types	Fix Mean Wage Tax
<i>Policy Rates:</i>						
Consumption Tax	30.4%	34.8%	19.3%	27.3%	32.6%	18%
Avg. Labor Income Tax	0.0%	0.0%	0.0%	0.0%	0.0%	13.4
Labor Tax Progressivity (τ_l)	0.35	0.36	0.28	0.37	0.36	0.35
Capital Income Tax	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
Wealth Tax	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
<i>Aggregate Quantities:</i>						
Wealth	16.6	18.2	22.4	11.5	16.0	1.9
Quality-Adj. Wealth	27.0	18.2	27.2	19.4	26.5	17.2
Hours	-12.1	-12.6	-12.4	-14.1	-11.9	-14.0
Output	0.4	-3.0	3.7	-3.3	0.5	-3.6
Consumption	-2.9	-10.3	1.2	-7.5	-4.5	-5.9
Investment	18.9	17.4	24.4	12.8	16.5	6.4
<i>Productivity:</i>						
TFP	3.1	0	1.4	2.5	3.2	5.2
Entrepreneurial Rate (Δ p.p.)	-2.9	n/a	-1	-1.2	-1.1	-5.0
<i>Prices:</i>						
Price of Capital	-20.9	-17.9	-18.5	-19.0	-20.6	-17.8
Wages	14.1	11.8	12.2	12.6	13.9	11.6
Borrowing Rate (Δ p.p.)	-1.0	-1.8		-1.3	-0.6	-0.6
<i>Inequality:</i>						
Wealth Gini	-2.2	-12.2	-1.7	-3.7	-1.1	-0.8
Earnings Gini	0.1	-0.2	0.1	0.0	0.1	0.3
Consumption Gini	-10.7	-21.5	-2.1	-12.9	-12.2	-7.6
<i>Welfare:</i>						
CE Welfare Gain	18.0	15.6	8.6	15.4	17.4	12.7
<i>Level</i>	4.0	-2.9	6.2	0.3	2.7	2.0
<i>Redistribution</i>	13.5	19.0	2.2	15	14.4	10.4
Pct. of Newborns Better Off	89	89	89	88	89	88

Note: In this table, we report the optimal policies (for the ex ante welfare of a newborn), where aggregate tax revenue is constrained to be non-negative for each tax instrument. See Table 5.

G Transition Path, Optimal Policy and Popular Consensus

An alternative approach to the transition path is imposing that any policy for the long run has to favor the *Median Voter* among the cohorts born before the policy announcement. Given that the second partial reform favors more than 60% of the cohort alive (voters) and

Table G.1 – Short-Run Welfare Effects of Fiscal Reforms: Pareto vs. Median Voter

Variable	Optimal Policy (Reform 3)	Optimal Policy (Pareto)	Optimal Policy (Median Voter)
Individuals on Impact			
Newborn	11.0	6.9	10.5
Age 22-30	8.4	6.0	9.0
Age 31-40	4.2	3.6	6.0
Age 41-50	-0.6	1.2	2.1
Age 51-64	-8.5	0.6	-4.7
Age 65+	-17.5	0.4	-13.1
<i>Average</i>	-3.6	2.3	-0.8
<i>Pct. Support</i>	40.1	100	50.9
Future Newborns			
+10 years	12.4	6.7	12.5
+20 years	13.1	6.8	13.3
+30 years	13.4	6.9	13.6
<i>Long Run</i>	18.0	18.0	16.1

the optimal policy without LSRA only around 40%, there could exist an intermediate policy reform, with a higher level of wage tax progressivity but lower than its optimal value, such that more than half of individuals benefit from it. The aim of this exercise is to quantify the long-run optimal policy that would be voted in by the living cohorts. We find that the tax progressivity is around 0.28 that would command a simple majority of agents is around 0.28. Relative to the benchmark policy this implies a reduction of progressivity of around 20%. Interestingly, relative to the optimal policy, it implies similar level of taxation for the bottom half of the distribution (-21% vs -24%) but lower tax rates for the top 1% (46% vs 55%). Given the distortionary effects of progressive wage taxation, in equilibrium this policy implies a lower consumption tax rate (28.2% vs. 30.4%).

Relative to the optimal policy in the benchmark experiment, lower progressivity is beneficial for middle-aged workers and detrimental for newborns, who now enjoy lower than optimal social insurance. At the same time, lower consumption tax rates decreases less the losses of the elderly. In aggregate, welfare on impact is still negative, but most people are better off, so the policy would have popular consensus. Interestingly, this policy would retain around 90% of the long-run gains of the benchmark policy (i.e., 16.1% vs 18.0% in CEV), with around 60% of these gains coming from redistribution.